

Data Replication Service

Real-Time Synchronization

Issue 57
Date 2024-02-28



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

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Technologies Co., Ltd.

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

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

Email: support@huawei.com

Security Declaration

Vulnerability

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

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

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

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

Contents

| | |
|--|-----------|
| 1 Synchronization Overview..... | 1 |
| 2 Data Synchronization Topologies..... | 13 |
| 3 To the Cloud..... | 16 |
| 3.1 From MySQL to MySQL..... | 16 |
| 3.2 From MySQL to PostgreSQL..... | 37 |
| 3.3 From MySQL to GaussDB Distributed..... | 54 |
| 3.4 From MySQL to GaussDB Primary/Standby..... | 72 |
| 3.5 From MySQL to GaussDB(DWS)..... | 91 |
| 3.6 From MySQL to GaussDB(for MySQL)..... | 114 |
| 3.7 From MySQL to MariaDB..... | 135 |
| 3.8 From PostgreSQL to PostgreSQL..... | 156 |
| 3.9 From PostgreSQL to GaussDB(DWS)..... | 181 |
| 3.10 From PostgreSQL to GaussDB Primary/Standby..... | 200 |
| 3.11 From PostgreSQL to GaussDB Distributed..... | 221 |
| 3.12 From Oracle to MySQL..... | 243 |
| 3.13 From Oracle to GaussDB(for MySQL)..... | 267 |
| 3.14 From Oracle to GaussDB Primary/Standby..... | 288 |
| 3.15 From Oracle to GaussDB Distributed..... | 311 |
| 3.16 From Oracle to DDM..... | 334 |
| 3.17 From Oracle to GaussDB(DWS)..... | 355 |
| 3.18 From Oracle to PostgreSQL..... | 376 |
| 3.19 From DDM to MySQL..... | 397 |
| 3.20 From DDM to GaussDB(DWS)..... | 411 |
| 3.21 From DDM to DDM..... | 425 |
| 3.22 From DB2 for LUW to GaussDB Primary/Standby..... | 438 |
| 3.23 From DB2 for LUW to GaussDB Distributed..... | 453 |
| 3.24 From DB2 for LUW to GaussDB(DWS)..... | 469 |
| 3.25 From TiDB to GaussDB(for MySQL)..... | 483 |
| 3.26 From Microsoft SQL Server to GaussDB(DWS)..... | 499 |
| 3.27 From Microsoft SQL Server to GaussDB Primary/Standby..... | 515 |
| 3.28 From Microsoft SQL Server to GaussDB Distributed..... | 531 |
| 3.29 From Microsoft SQL Server to Microsoft SQL Server..... | 547 |

| | |
|---|-------------|
| 3.30 From MongoDB to DDS..... | 562 |
| 3.31 From MariaDB to MariaDB..... | 578 |
| 3.32 From MariaDB to MySQL..... | 598 |
| 3.33 From MariaDB to GaussDB(for MySQL)..... | 620 |
| 3.34 From GaussDB(for MySQL) to GaussDB(for MySQL)..... | 642 |
| 4 Out of the Cloud..... | 663 |
| 4.1 From MySQL to MySQL..... | 663 |
| 4.2 From MySQL to Kafka..... | 685 |
| 4.3 From MySQL to CSS/ES..... | 702 |
| 4.4 From MySQL to Oracle..... | 719 |
| 4.5 From MySQL to MariaDB..... | 736 |
| 4.6 From DDM to MySQL..... | 757 |
| 4.7 From DDM to Oracle..... | 771 |
| 4.8 From DDM to Kafka..... | 786 |
| 4.9 From DDS to MongoDB..... | 798 |
| 4.10 From DDS to Kafka..... | 812 |
| 4.11 From PostgreSQL to PostgreSQL..... | 823 |
| 4.12 From PostgreSQL to Kafka..... | 844 |
| 4.13 From GaussDB Distributed to MySQL..... | 857 |
| 4.14 From GaussDB Distributed to Oracle..... | 878 |
| 4.15 From GaussDB Distributed to GaussDB(DWS)..... | 900 |
| 4.16 From GaussDB Distributed to Kafka..... | 920 |
| 4.17 From GaussDB Distributed to GaussDB Distributed..... | 935 |
| 4.18 From GaussDB Distributed to GaussDB Primary/Standby..... | 956 |
| 4.19 From GaussDB Primary/Standby to MySQL..... | 977 |
| 4.20 From GaussDB Primary/Standby to Oracle..... | 999 |
| 4.21 From GaussDB Primary/Standby to GaussDB(DWS)..... | 1022 |
| 4.22 From GaussDB Primary/Standby to Kafka..... | 1042 |
| 4.23 From GaussDB Primary/Standby to GaussDB Distributed..... | 1058 |
| 4.24 From GaussDB Primary/Standby to GaussDB Primary/Standby..... | 1078 |
| 4.25 From GaussDB(for MySQL) to MySQL..... | 1098 |
| 4.26 From GaussDB(for MySQL) to GaussDB(DWS)..... | 1115 |
| 4.27 From GaussDB(for MySQL) to Kafka..... | 1134 |
| 4.28 From GaussDB(for MySQL) to CSS/ES..... | 1149 |
| 4.29 From GaussDB(for MySQL) to Oracle..... | 1166 |
| 4.30 From MariaDB to MariaDB..... | 1181 |
| 4.31 From Microsoft SQL Server to Kafka..... | 1200 |
| 5 Between Self-built Databases..... | 1212 |
| 5.1 From MySQL to Kafka..... | 1212 |
| 5.2 From MySQL to CSS/ES..... | 1229 |
| 5.3 From MySQL to GaussDB Primary/Standby..... | 1246 |
| 5.4 From MySQL to GaussDB Distributed..... | 1264 |

| | |
|---|-------------|
| 5.5 From Oracle to Kafka..... | 1283 |
| 5.6 From Oracle to GaussDB Primary/Standby..... | 1302 |
| 5.7 From Oracle to GaussDB Distributed..... | 1324 |
| 5.8 From PostgreSQL to Kafka..... | 1347 |
| 5.9 From GaussDB Primary/Standby to MySQL..... | 1360 |
| 5.10 From GaussDB Primary/Standby to Oracle..... | 1381 |
| 5.11 From GaussDB Primary/Standby to Kafka..... | 1403 |
| 5.12 From GaussDB Primary/Standby to GaussDB Distributed..... | 1419 |
| 5.13 From GaussDB Primary/Standby to GaussDB Primary/Standby..... | 1439 |
| 5.14 From GaussDB Distributed to MySQL..... | 1459 |
| 5.15 From GaussDB Distributed to Oracle..... | 1480 |
| 5.16 From GaussDB Distributed to Kafka..... | 1502 |
| 5.17 From GaussDB Distributed to GaussDB Distributed..... | 1518 |
| 5.18 From GaussDB Distributed to GaussDB Primary/Standby..... | 1539 |
| 5.19 From DB2 for LUW to GaussDB Primary/Standby..... | 1560 |
| 5.20 From DB2 for LUW to GaussDB Distributed..... | 1575 |
| 5.21 From Microsoft SQL Server to Kafka..... | 1590 |
| 6 Task Management..... | 1602 |
| 6.1 Creating a Synchronization Task..... | 1602 |
| 6.2 Querying the Synchronization Progress..... | 1618 |
| 6.3 Viewing Synchronization Logs..... | 1621 |
| 6.4 Comparing Synchronization Items..... | 1622 |
| 6.5 Managing Objects..... | 1639 |
| 6.5.1 Editing Synchronization Objects..... | 1640 |
| 6.5.2 Importing Synchronization Objects..... | 1643 |
| 6.5.3 Mapping Object Names..... | 1644 |
| 6.5.4 Viewing Synchronization Mapping Information..... | 1650 |
| 6.5.5 Processing Data..... | 1651 |
| 6.6 Managing Parameters..... | 1662 |
| 6.6.1 Changing Task Parameters..... | 1663 |
| 6.6.2 Viewing Parameter Change History..... | 1666 |
| 6.7 Task Life Cycle..... | 1666 |
| 6.7.1 Viewing Task Details..... | 1666 |
| 6.7.2 Modifying Task Information..... | 1667 |
| 6.7.3 Modifying Connection Information..... | 1668 |
| 6.7.4 Modifying the Flow Control Mode..... | 1669 |
| 6.7.5 Editing a Synchronization Task..... | 1673 |
| 6.7.6 Resuming a Synchronization Task..... | 1673 |
| 6.7.7 Skipping Data and Resuming a Synchronization Task..... | 1674 |
| 6.7.8 Pausing a Synchronization Task..... | 1675 |
| 6.7.9 Resetting a Synchronization Task..... | 1678 |
| 6.7.10 Resetting Synchronization Position..... | 1681 |

| | |
|---|-------------|
| 6.7.11 Restarting a Synchronization Task..... | 1682 |
| 6.7.12 Skipping DDL Statements..... | 1683 |
| 6.7.13 Performing a Switchover for a Dual-AZ Task..... | 1683 |
| 6.7.14 Cloning a Synchronization Task..... | 1685 |
| 6.7.15 Changing Specifications..... | 1687 |
| 6.7.16 Unsubscribing from a Yearly/Monthly Task..... | 1688 |
| 6.7.17 Stopping a Synchronization Task..... | 1690 |
| 6.7.18 Deleting a Synchronization Task..... | 1692 |
| 6.7.19 Importing Tasks in Batches..... | 1692 |
| 6.7.20 Task Statuses..... | 1694 |
| 7 Tag Management..... | 1697 |
| 8 Connection Diagnosis..... | 1700 |
| 9 Diagnosis of Abnormal Records..... | 1702 |
| 10 Interconnecting with CTS..... | 1703 |
| 10.1 Key Operations Recorded by CTS..... | 1703 |
| 10.2 Viewing Traces..... | 1703 |
| 11 Interconnecting with Cloud Eye..... | 1705 |
| 11.1 Supported Metrics..... | 1705 |
| 11.2 Configuring Alarm Rules..... | 1710 |
| 11.3 Viewing Monitoring Metrics..... | 1712 |
| 12 Interconnecting with LTS..... | 1714 |
| 12.1 Log Reporting..... | 1714 |
| 12.2 Viewing and Downloading Logs..... | 1715 |
| 13 Operation Reference in Synchronization Scenarios..... | 1718 |
| 13.1 Kafka Message Format..... | 1718 |
| 13.2 Kafka Authentication..... | 1727 |
| 13.3 Forcibly Stopping Synchronization of GaussDB Distributed..... | 1730 |
| 13.4 Forcibly Stopping Synchronization of GaussDB Primary/Standby..... | 1731 |
| 13.5 Forcibly Stopping Synchronization of PostgreSQL..... | 1732 |
| 13.6 Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL..... | 1734 |
| 14 Appendix..... | 1738 |
| 14.1 General Usage Suggestions..... | 1738 |
| 14.2 DRS Pre-Check Items..... | 1739 |
| A Change History..... | 1759 |

1 Synchronization Overview

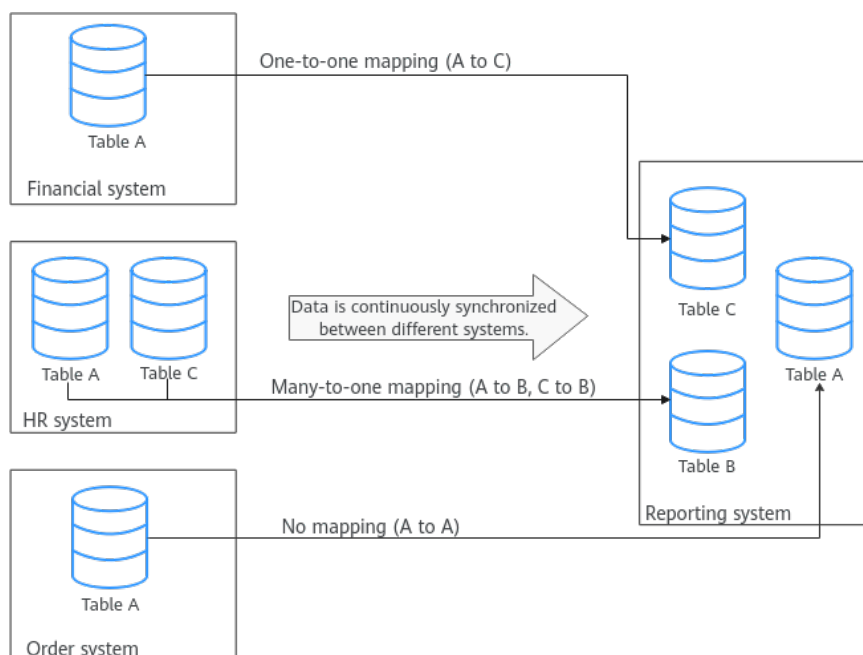
Real-time synchronization refers to the real-time flow of key service data from sources to destinations while consistency of data can be ensured.

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

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

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

Figure 1-1 Many-to-one real-time synchronization process



Supported Database Types

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

Table 1-1 Synchronization scheme

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

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

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

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

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

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

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

| Source DB Type | Destination DB Type | Synchronization Mode | Related Documents |
|----------------|--|---|---|
| | <ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL | 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 |
| | <ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS | Full+Incremental | From GaussDB(for MySQL) to Oracle |
| DDS | <ul style="list-style-type: none"> Self-built MongoDB MongoDB on other clouds | Incremental synchronization supports the following modes: Replica set -> Replica set Cluster -> Cluster (the source cluster version must be 4.0 or later) | From DDS to MongoDB |
| | Kafka | Incremental | From DDS to Kafka |

| Source DB Type | Destination DB Type | Synchronization Mode | Related Documents |
|---|--|---|---|
| DB2 for LUW | GaussDB distributed | Full Full+Incremental | From DB2 for LUW to GaussDB Primary/Standby (To the cloud) From DB2 for LUW to GaussDB Primary/Standby (self-built - self-built) |
| | GaussDB primary/standby | Full Full+Incremental | From DB2 for LUW to GaussDB Distributed (To the cloud) From DB2 for LUW to GaussDB Distributed (self-built - self-built) |
| | GaussDB(DWS) | Full Full+Incremental | From DB2 for LUW to GaussDB(DWS) |
| TiDB | GaussDB(for MySQL) | Full+Incremental | From TiDB to GaussDB(for MySQL) |
| <ul style="list-style-type: none"> • On-premises Microsoft SQL Server database • Microsoft SQL Server databases on an ECS • Microsoft SQL Server databases on other clouds • RDS for SQL Server | GaussDB(DWS) | Full+Incremental | From Microsoft SQL Server to GaussDB(DWS) |
| | GaussDB primary/standby | Incremental Full Full+Incremental | From Microsoft SQL Server to GaussDB Primary/Standby |
| | GaussDB distributed | Incremental Full Full+Incremental | From Microsoft SQL Server to GaussDB Distributed |
| | Microsoft SQL Server NOTE Only whitelisted users can use this function. | Full+Incremental | From Microsoft SQL Server to Microsoft SQL Server |

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

| Source DB Type | Destination DB Type | Synchronization Mode | Related Documents |
|---|---|---------------------------------|--|
| <ul style="list-style-type: none"> On-premises MariaDB databases MariaDB databases on an ECS MariaDB databases on other clouds | RDS for MySQL NOTE Only whitelisted users can use this function. | Incremental Full+Incremental | From MariaDB to MySQL |
| | GaussDB(for MySQL) NOTE Only whitelisted users can use this function. | Incremental Full+Incremental | From MariaDB to GaussDB(for MySQL) |
| RDS for MariaDB | <ul style="list-style-type: none"> On-premises MariaDB databases MariaDB databases on an ECS MariaDB databases on other clouds | Full+Incremental | From MariaDB to MariaDB |


2 Data Synchronization Topologies

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

 **NOTE**


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

One-to-One Real-Time Synchronization




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

One-to-Many Real-Time Synchronization



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

Many-to-One Real-Time Synchronization



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

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

3 To the Cloud

3.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 3-1 Supported databases

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

Database Account Permission Requirements

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

Table 3-2 Database account permission

| Type | Full Synchronization | Full+Incremental Synchronization and Incremental Synchronization |
|---------------------------|---|---|
| Source database user | The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT | The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT |
| Destination database user | The user must have the following minimum permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES The root account of the RDS for MySQL DB instance has the preceding permissions by default. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. | |

 **NOTE**

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

Supported Synchronization Objects

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

Table 3-3 Supported synchronization objects

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

Precautions

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

Table 3-4 Precautions

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

| Type | Restrictions |
|------------------|---|
| Use restrictions | <p>General</p> <ul style="list-style-type: none"> ● You are not advised to modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. ● Do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● Do not clear binlogs on the source database. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. <p>Full synchronization phase</p> <ul style="list-style-type: none"> ● When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. ● Modifying MyISAM tables may cause data inconsistency. ● During full data synchronization, a lot of binlogs are generated in the destination database, occupying too much storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. <p>Incremental synchronization phase</p> <ul style="list-style-type: none"> ● If the session variable character_set_client is set to binary, some data may include garbled characters. ● During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? ● You can add additional synchronization objects. ● Resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down. ● Some DDL statements are supported. <ul style="list-style-type: none"> – In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, |

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

| Type | Restrictions |
|--------------------|---|
| Other restrictions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see Syntax Differences Between MySQL Versions. ● During table name mapping, tables on which views, stored procedures, and functions depend cannot be synchronized, and foreign key constraints of tables cannot be synchronized. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the sources and destinations are RDS instances, database mapping is required. ● If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. ● Braces ({} in a database or table name in the source database must be used in pairs. Otherwise, the table structure synchronization may fail. After data mapping, if braces ({} in a database or name mapped to the destination database are not in pairs, the table structure may fail to be synchronized. ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● The partitioned table does not support column mapping. ● If the source and destination sides are RDS for MySQL instances, transparent data encryption (TDE) is not supported, and tables with the encryption function cannot be created. ● If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● After a synchronization task is created, the destination database cannot be set to read-only. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination database cannot be null or have default values. If newly inserted |

| Type | Restrictions |
|------|---|
| | <p>data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. • If you create a many-to-one synchronization task, refer to Constraints and Operation Suggestions on Many-to-One Scenario. • For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail. • The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. |

Prerequisites

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

Procedure

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

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

Figure 3-1 Synchronization task information

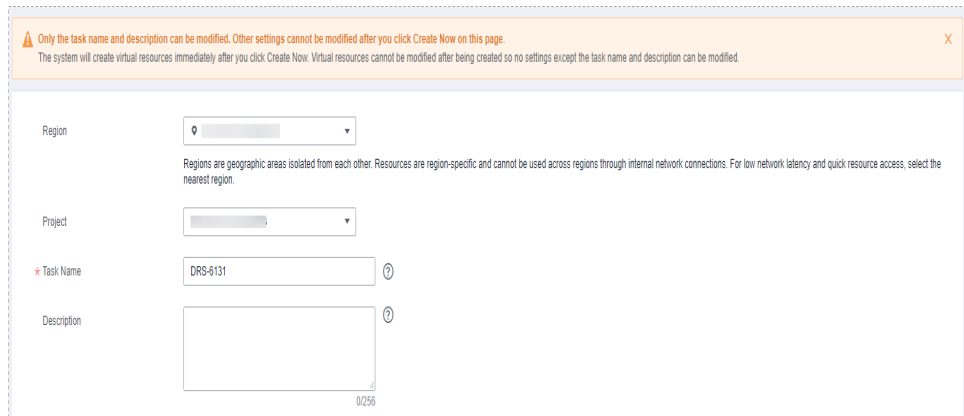


Table 3-5 Task information

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

- Synchronization instance information

Figure 3-2 Synchronization instance information

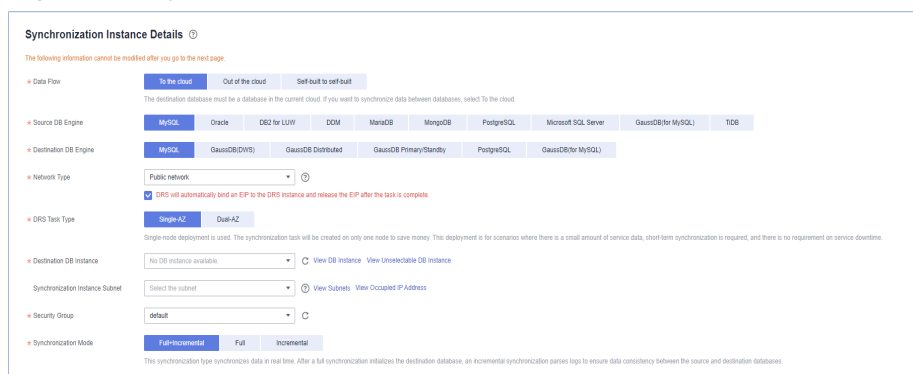


Table 3-6 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|--|
| Destination DB Instance | <p>The RDS DB instance you created.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The destination DB instance cannot be a read replica. - The source and destination DB instances can be the same DB instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE</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"> - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 3-3 Task type



Table 3-7 Task type information

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

- Enterprise Project and Tags

Figure 3-5 Enterprise projects and tags

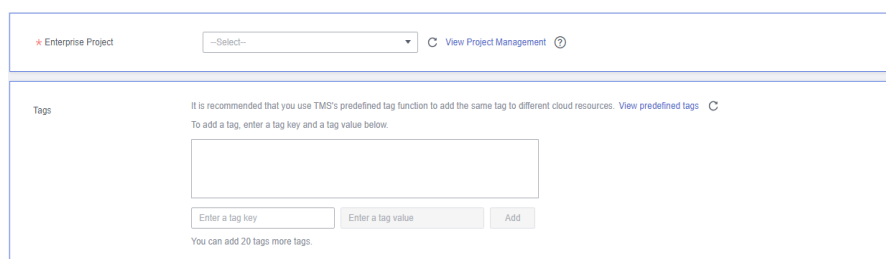


Table 3-8 Enterprise Project and Tags

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

 **NOTE**

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

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

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

- Source database information

Figure 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 3-9 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |
| 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 disabled, your data may be at risk. |

NOTE

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


- Destination database information

Figure 3-7 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-10 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 disabled, your data may be at risk. |

 **NOTE**

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

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

Figure 3-8 Synchronization mode

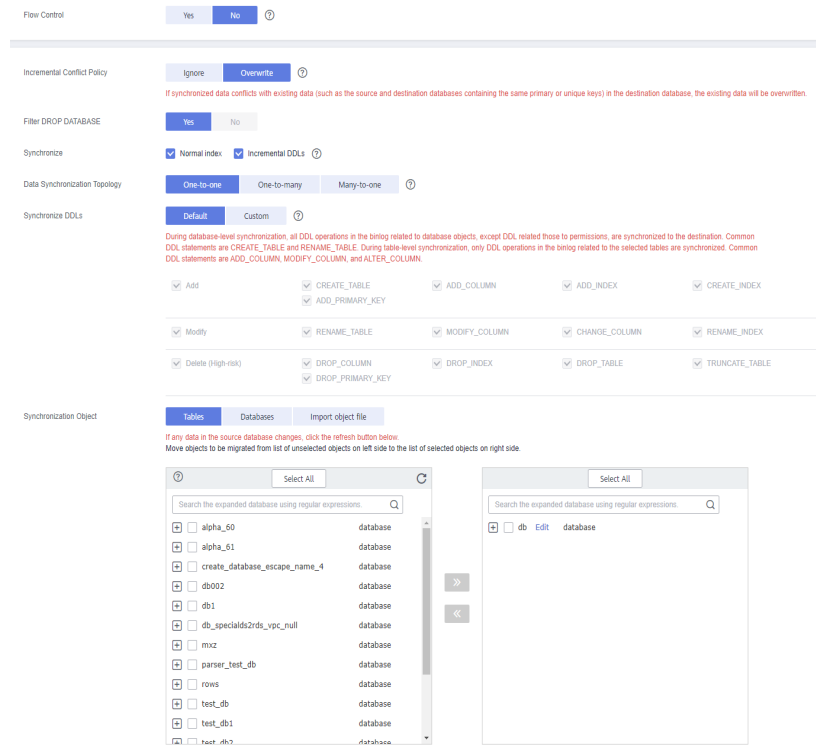
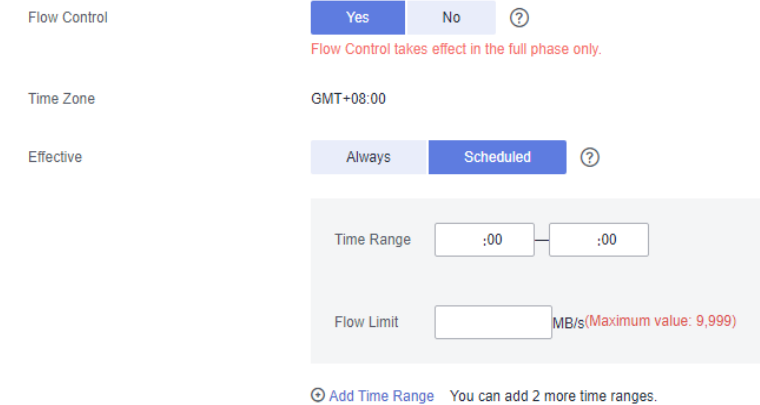



Table 3-11 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize 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 position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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> |

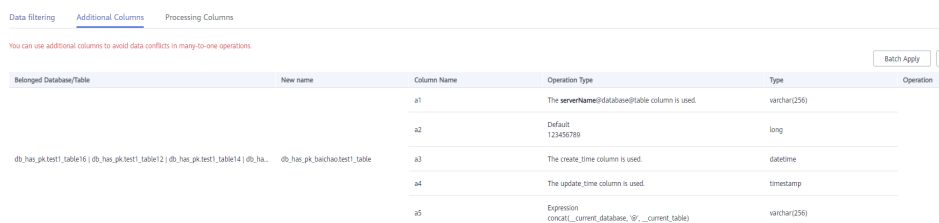
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: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

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

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

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

Figure 3-10 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_', current_table) | varchar(256) | |

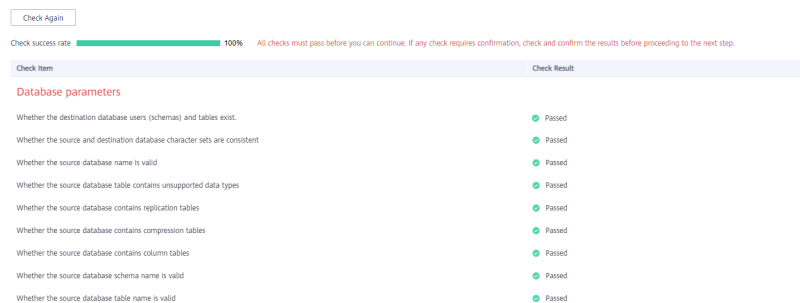
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Solutions to Failed Check Items](#) 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-12 Task startup settings

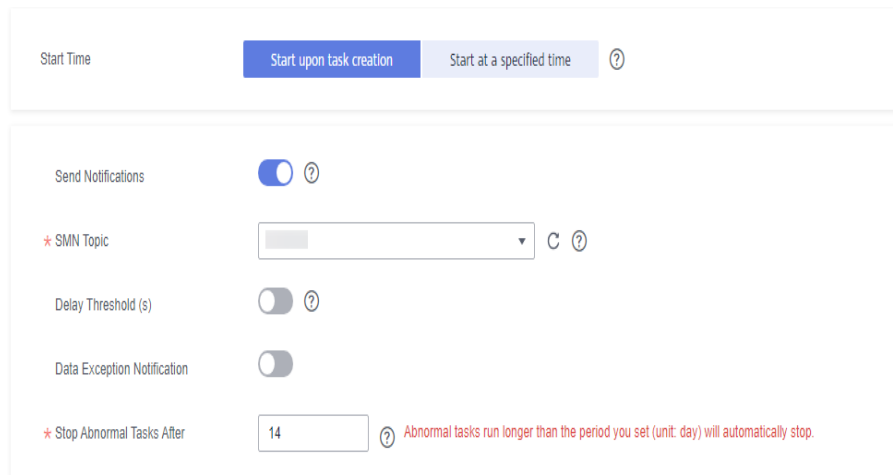



Table 3-12 Task startup settings

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

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

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

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

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

3.2 From MySQL to PostgreSQL

Supported Source and Destination Databases

Table 3-13 Supported databases

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

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-14](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-14 Database account permission

| Type | Full Synchronization | Full+Incremental Synchronization |
|---------------------------|--|---|
| Source database user | The user must have the following minimum permissions: SELECT | The user must have the following minimum permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT |
| Destination database user | The account of the RDS for PostgreSQL instance has the permissions by default. | |

 NOTE

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

Supported Synchronization Objects

[Table 3-15](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-15 Supported synchronization objects

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

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-16](#).

Table 3-16 Precautions

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

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

| Type | Restrictions |
|------|---|
| | <p data-bbox="587 297 799 331">Stopping a task</p> <ul data-bbox="587 340 1422 555" style="list-style-type: none"><li data-bbox="587 340 1422 477">• Stop a task normally. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the synchronization status reported by the synchronization log.<li data-bbox="587 486 1422 555">• Forcibly stop a task. If you forcibly stop a task, DRS resources will be released. <p data-bbox="587 564 810 598">Troubleshooting</p> <ul data-bbox="587 607 1422 701" style="list-style-type: none"><li data-bbox="587 607 1422 701">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting. |

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

| Type | Restrictions |
|------|--|
| | <p>binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.</p> <ul style="list-style-type: none"> - If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. • The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination database cannot be null or have default values. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will become null, which does not meet the requirements of the destination database and will cause the task to fail. - Assume that extra columns on the destination database must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source database to the destination database, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination database and will cause the task to fail. • Only data that violates the non-null constraint and data of the char or varchar type that exceeds the field length limit can be recorded. • If tables without primary keys contain large fields (BLOB, TEXT, CLOB, NCLOB, or BYTEA), data of the large fields may be inconsistent during incremental synchronization. • If the source table to be synchronized has the AUTO_INCREMENT attribute, DRS automatically updates the start value of the PostgreSQL auto-increment sequence corresponding to the integer sequence of the table when the task is complete. The updated value is the maximum value of the sequence plus 10,000. |

Procedure

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

Figure 3-13 Synchronization task information

Table 3-17 Task information

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

- Synchronization instance details

Figure 3-14 Synchronization instance details

Table 3-18 Synchronization instance settings

| Parameter | Description |
|-----------|------------------------------|
| Data Flow | Select To the cloud . |

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

- Task type

Figure 3-15 Task type



Table 3-19 Task type information

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

- Enterprise Project and Tags

Figure 3-16 Enterprise projects and tags

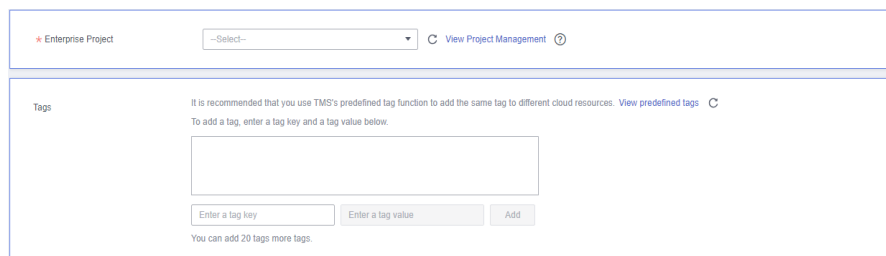


Table 3-20 Enterprise Project and Tags

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

 **NOTE**

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

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

Figure 3-17 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

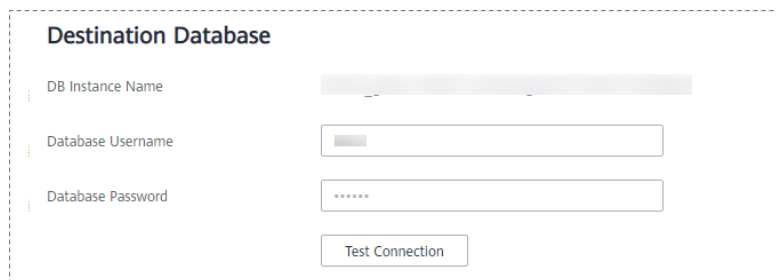
Table 3-21 Self-build on ECS - source database information

| 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 disabled, your data may be at risk. |

 **NOTE**

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

Figure 3-18 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 3-22 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

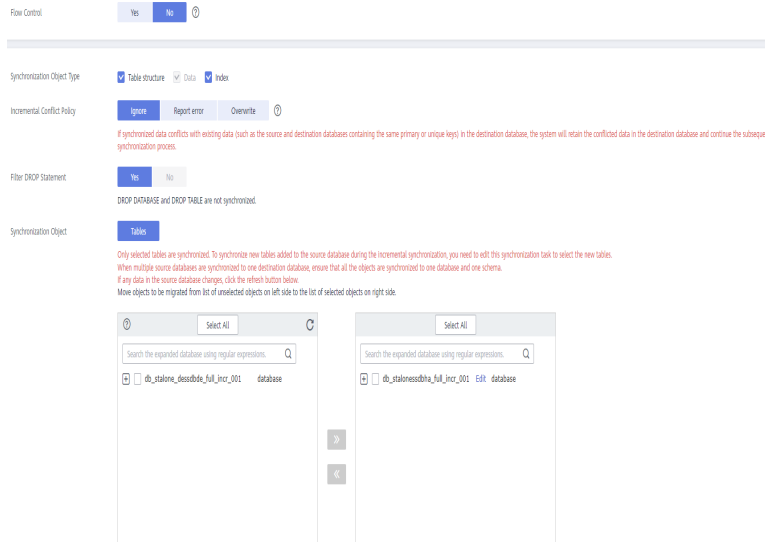
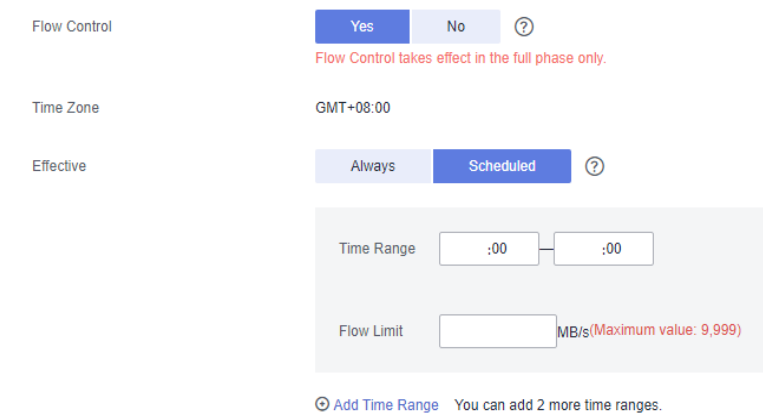



Table 3-23 Synchronization Object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization. |

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

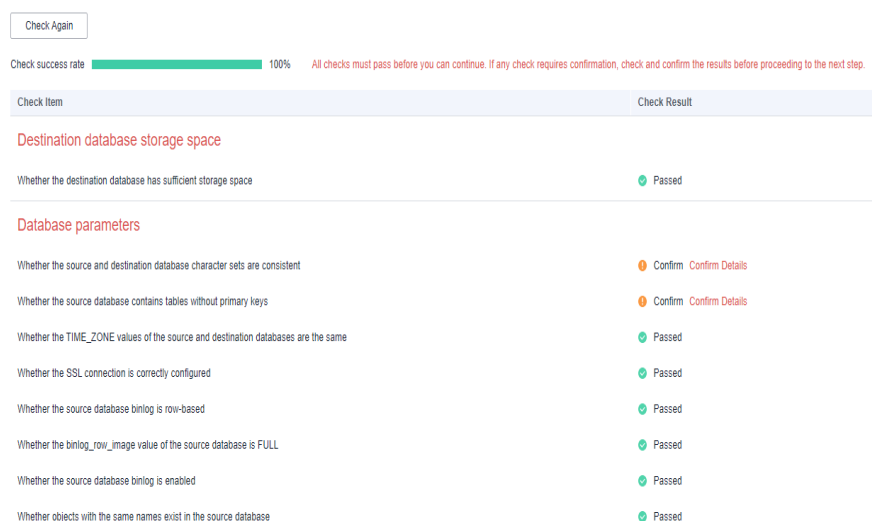
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-22 Task startup settings

Table 3-24 Task startup settings

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

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

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

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

----End

3.3 From MySQL to GaussDB Distributed

Supported Source and Destination Databases

Table 3-25 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 |

Suggestions

⚠ CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-26 Precautions

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

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● Only table structures, table data, and indexes 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. ● Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed. ● The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and interval. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● During 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 a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. – If the source database is 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Restrictions |
|------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • The 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 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 table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • 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 an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated. • 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. |

| 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. • Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. • If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. • Floating-point data is approximate numbers and depends on the OS platform and underlying implementation. FLOAT and DOUBLE data is inaccurate. If you synchronize floating-point data between MySQL and GaussDB, there may be some differences. For details, see the official MySQL documentation. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. • Two-phase commit 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

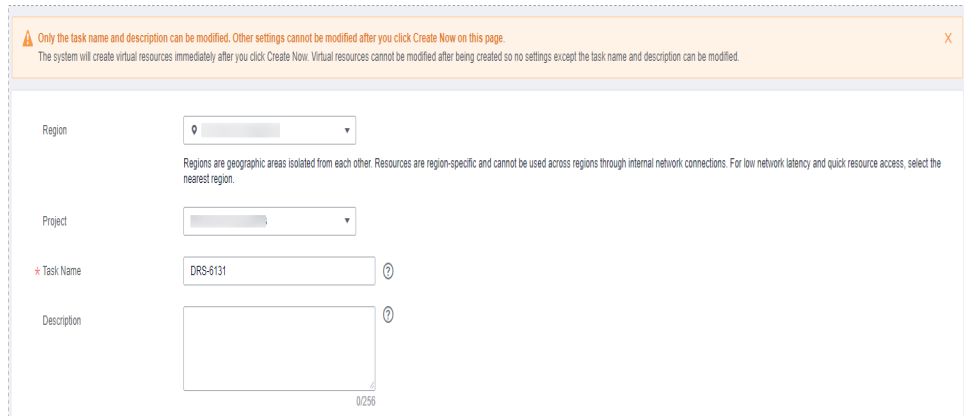


Table 3-27 Task information

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

- Synchronization instance details

Figure 3-24 Synchronization instance details

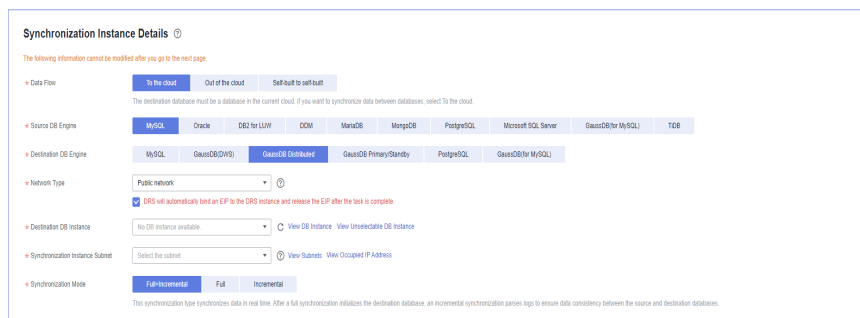


Table 3-28 Synchronization instance settings

| Parameter | Description |
|-----------|------------------------------|
| Data Flow | Select To the cloud . |

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

| Parameter | Description |
|----------------------|--|
| Synchronization mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE</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. Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 3-25 Task type




Table 3-29 Task type information

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

- Enterprise Project and Tags

Figure 3-26 Enterprise projects and tags

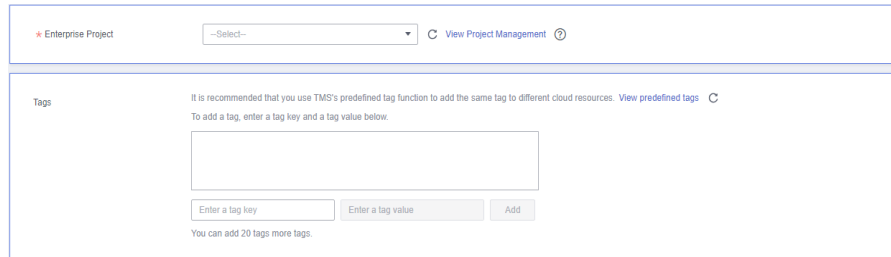


Table 3-30 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-27 Source database information

Source Database

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

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 3-31 Source database settings

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

 **NOTE**

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

Figure 3-28 Destination database information

The screenshot shows a configuration form titled "Destination Database". It contains 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 3-32 Destination database settings

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

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

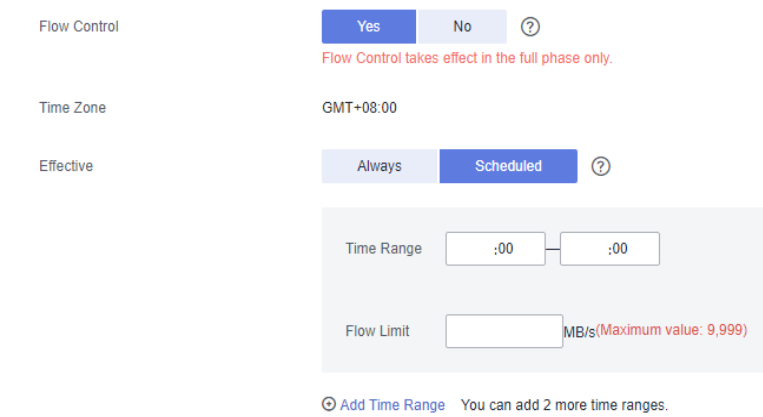
Figure 3-29 Synchronization mode


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

- Flow Control:** Radio buttons for "Yes" and "No".
- Synchronization Object Type:** Radio buttons for "Table structure", "Data", and "Index". A note below states: "When you manually create a table structure in the destination database, for details about the data type, see Mapping Data Types."
- Incremental Conflict Policy:** Radio buttons for "Ignore", "Report error", and "Overwrite". A note below states: "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."
- Synchronization Object:** Radio buttons for "Tables" and "Import object file". A note below 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 databases with checkboxes, and the right panel shows a search bar and a list of selected objects with "Edit" links.

Table 3-33 Synchronization mode and object

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

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

- Step 5** On the **Process Data** page, set the filtering rules for data processing.
- If data processing is not required, click **Next**.
 - If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-31 Processing data

| Selected Database/Table | Filter name | Column Name | Operation Type | Type | Operation |
|-------------------------|-------------|-------------|----------------|------|-----------|
| db1.pak_011 | db1.pak_011 | -- | -- | -- | Add |
| db1.test2 | db1.test2 | -- | -- | -- | Add |
| db1.test3 | db1.test3 | -- | -- | -- | Add |

- Step 6** On the **Check Task** page, check the synchronization task.
- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
 - If all check items are successful, click **Next**.

Figure 3-32 Pre-check

| Check Item | Check Result |
|---|-------------------------|
| Database parameters | |
| Whether there are source database tables that do not contain primary keys | Confirm Confirm Details |
| Whether the character set type is supported | Passed |
| Whether the destination database contains the configured databases | Passed |
| Whether the SSL connection is correctly configured | Passed |
| Whether the database names mapped to the destination database are valid | Passed |
| Whether the source database binlog is row-based | Passed |

NOTE

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

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

Figure 3-33 Task startup settings





Table 3-34 Task startup settings

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

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

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

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

----End

3.4 From MySQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-35 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 |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-36 Precautions

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

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● Only table structures, table data, and indexes can be synchronized. Other database objects such as stored procedures cannot be synchronized. ● Incremental synchronization does not support synchronization of DDL. ● Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed. ● The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and interval. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● During 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 a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. – If the source database is 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Restrictions |
|------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • The 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 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 table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • 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 an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated. • 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. |

| 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. • Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. • If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. • 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. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. • If 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 3-34 Synchronization task information

Table 3-37 Task information

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

- Synchronization instance details

Figure 3-35 Synchronization instance details

Table 3-38 Synchronization instance settings

| Parameter | Description |
|-----------|------------------------------|
| Data Flow | Select To the cloud . |

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

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

- Task type

Figure 3-36 Task type



Table 3-39 Task type information

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

- Enterprise Project and Tags

Figure 3-37 Enterprise projects and tags




Table 3-40 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-38 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 3-41 Source database settings

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

 **NOTE**

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

Figure 3-39 Destination database information

Destination Database

DB Instance Name:

Database Username:

Database Password:

Table 3-42 Destination database settings

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

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

Figure 3-40 Synchronization mode

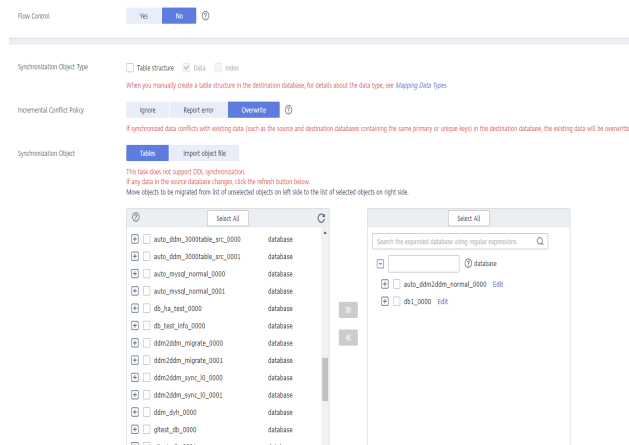
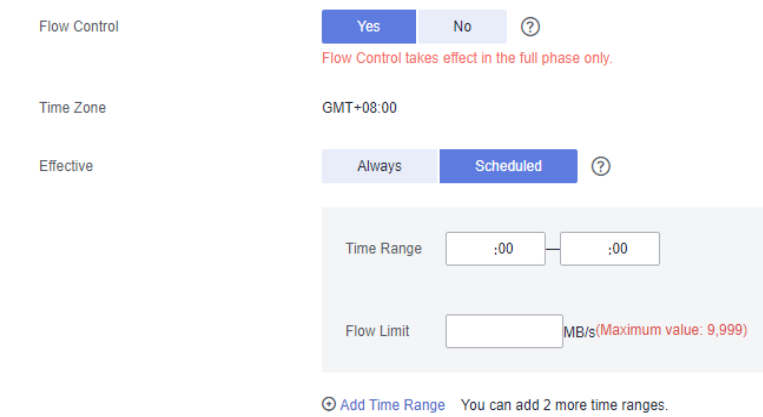



Table 3-43 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-41 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

- Step 5** On the **Process Data** page, set the filtering rules for data processing.
- If data processing is not required, click **Next**.
 - If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 3-42 Processing data

| Belonged Database/Table | New name | Column Name | Operation Type | Type | Operation |
|-------------------------|--------------|-------------|----------------|------|-----------|
| db1_jack_011 | db1_jack_011 | -- | -- | -- | Add |
| db1_test2 | db1_test2 | -- | -- | -- | Add |
| db1_test3 | db1_test3 | -- | -- | -- | Add |

- Step 6** On the **Check Task** page, check the synchronization task.
- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
 - If all check items are successful, click **Next**.

Figure 3-43 Pre-check

| Check Item | Check Result |
|---|-------------------------|
| Destination database storage space | |
| Whether the destination database has sufficient storage space | Passed |
| Database parameters | |
| Whether the source and destination database character sets are consistent | Confirm Confirm Details |
| Whether the source database contains tables without primary keys | Confirm Confirm Details |
| Whether the TIME_ZONE values of the source and destination databases are the same | Passed |
| Whether the SSL connection is correctly configured | Passed |
| Whether the source database binlog is row-based | Passed |
| Whether the binlog_row_image value of the source database is FULL | Passed |
| Whether the source database binlog is enabled | Passed |
| Whether objects with the same names exist in the source database | Passed |

NOTE

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

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

Figure 3-44 Task startup settings





Table 3-44 Task startup settings

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

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

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

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

----End

3.5 From MySQL to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-45 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(DWS) cluster 8.1.3 and 8.2.0 |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If 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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-46 Precautions

| Type | Restrictions |
|----------------------|---|
| Database permissions | <ul style="list-style-type: none">• The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT.• The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE. |

| Type | Restrictions |
|------------------------|--|
| Synchronization object | <ul style="list-style-type: none"> ● Only tables, common indexes (B-Tree indexes), and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, and unique indexes cannot be synchronized. ● Comment is supported in full synchronization mode, but not supported in incremental synchronization mode. ● The following data types are not supported: XML, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● 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. ● The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*:80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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 a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. – If the source database is 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 <'>.`\ " |
| 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. If the time zone settings are inconsistent, the values related to the time data type are inconsistent after the synchronization. After the time zones are changed to the same, the values are consistent. |

| 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. ● If there is a unique key when a table without a primary key is synchronized, data may fail to be written due to data conflicts. ● Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. ● If tables without primary keys contain large fields (BLOB, TEXT, CLOB, NCLOB, or BYTEA), data of the large fields may be inconsistent during incremental synchronization. ● If the source 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● MySQL 5.7 filters CHECK constraints. If the source database version is 5.7, you are not advised to use CHECK constraints, which may cause data inconsistencies. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> ● 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. ● 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 MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB(DWS), the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: hash value + original constraint name (which may be truncated) + _key, or table name_original index name. ● In the full synchronization phase, only B-Tree indexes are synchronized. Other indexes are not synchronized by default. If there are too many GaussDB(DWS) indexes, the storage space and data import performance will be affected. You are advised to create indexes based on service requirements. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> ● In a full synchronization for the table structure, the length of char, varchar, nvarchar, enum, and set characters in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. During incremental partition table synchronization, partition tables are created based on the types supported by GaussDB(DWS), which may fail. ● When the timestamp data is to be synchronized in full synchronization, the on update current_timestamp syntax in the default value will not be synchronized to the destination GaussDB(DWS). ● 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, DROP_CONSTRAINT, ADD_CONSTRAINT, CREATE_INDEX, DROP_INDEX, RENAME_INDEX, and RENAME_COLUMN. You can select the DDL operations to be synchronized on the object selection page as required. - If you rename a column in many-to-one synchronization, you must stop services. Otherwise, data inconsistency may occur. - In many-to-one scenarios, you are advised to synchronize ADD_COLUMN only. Other DDL synchronization may cause task failures or data inconsistency due to destination table changes. - In many-to-one scenarios, when running ADD_COLUMN, ensure that the types of columns added to each table are the same. Otherwise, the task may fail. - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. - When an index is creating using a DDL statement, if the table name and index name in the statement contain more than 63 characters, duplicate names may occur. As a result, the index fails to be created. - If a primary key is added to a table that does not have a primary key in the source database, the DDL operation must contain the first column. Otherwise, the task fails. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> - When a DDL operation is performed in the incremental phase, if the destination table is not found, the DDL operation will be ignored. - In the incremental phase, if CHANGE COLUMN is performed in the source database to modify a column and the column is a distribution column in destination GaussDB(DWS), the statement will be ignored because GaussDB(DWS) does not support distribution column modification. - In the incremental phase, the RENAME INDEX operation is not supported because the index rules of the source and destination databases are different. - In the incremental phase, you are not advised to combine CHAR(0) with other characters for data synchronization. For example, inserting CONCAT('a',CHAR(0),'b') may cause data inconsistencies. - In the incremental phase, you are not advised to combine CHAR(34) (double quotation marks) and CHAR(92) (backslash) with other characters for JSON data synchronization. For example, inserting JSON_OBJECT('\$\$.0', CONCAT('a', CHAR(34), 'b')) may cause data inconsistencies because an extra backslash (\) will be added after it is synchronized to the destination end. • The SQL statements in the DDL for incremental synchronization come from the native SQL statements executed by users. The table structure syntax for full synchronization is the SQL statements processed by the SQL syntax layer. The types or constraints of these SQL statements are different. As a result, the table structures or field types created in full synchronization and incremental synchronization may be different. For details, see Suggestions on Synchronizing Data to GaussDB(DWS). • During incremental synchronization, DDL operations (such as ALTER TABLE ddl_test ADD COLUMN c2 AFTER/FIRST c1) for adding columns to a specified position are not supported. DRS will delete the AFTER/FIRST attribute, which may cause column sequence inconsistency. • 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. For details, see GaussDB(DWS) Parameter Description. Integer types: TINYINT, SMALLINT, INT, BIGINT, and NUMERIC/DECIMAL |

| Type | Restrictions |
|------|---|
| | <p>Character types: CHAR, BPCHAR, VARCHAR, VARCHAR2, and NVARCHAR2</p> <p>Date/time types: DATE, TIME, TIMETZ, TIMESTAMP, TIMESTAMPTZ, INTERVAL, and SMALLDATETIME</p> <p>For a table without a primary key, a column that can be used as a distribution column is specified as the distribution column. If all column types cannot be used as distribution columns, the table fails to be created and DRS synchronization is interrupted.</p> <ul style="list-style-type: none"> • During an incremental synchronization, tables whose primary key type is binary, text, blob, or clob cannot be deleted or updated. • 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 3-45 Synchronization task information

Table 3-47 Task information

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

- Synchronization instance details

Figure 3-46 Synchronization instance details

Table 3-48 Synchronization instance settings

| Parameter | Description |
|-----------|------------------------------|
| Data Flow | Select To the cloud . |

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

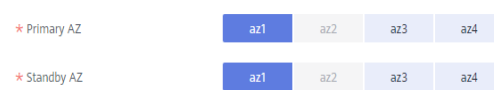
| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. 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. |

- Task Type

Figure 3-47 Task type



Table 3-49 Task type information

| Parameter | Description |
|----------------|--|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-48 AZ</p>  |

- Enterprise Project and Tags

Figure 3-49 Enterprise projects and tags

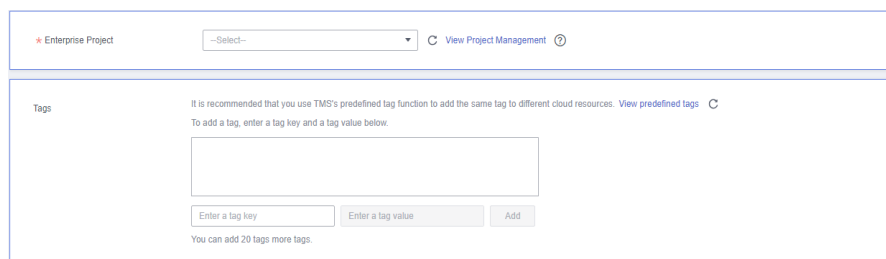


Table 3-50 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-50 Source database information

Source Database

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

Database Type: self-built database RDS DB Instance

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

Database Username:

Database Password:

SSL Connection:

Table 3-51 Source database settings

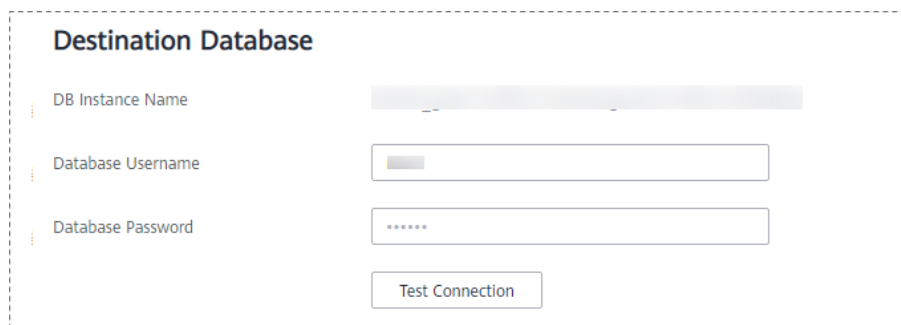
| Parameter | Description |
|-------------------|---|
| Database Type | Select Self-built database or RDS DB instance . |
| DB Instance Name | Select the RDS DB instance to be synchronized. |
| Database Username | The username for accessing the source database. |

| Parameter | Description |
|-------------------|---|
| 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 3-51 Destination database information



The screenshot shows a configuration form titled "Destination Database". It contains three input fields: "DB Instance Name" (with a dropdown arrow), "Database Username" (with a text input field), and "Database Password" (with a masked input field showing six asterisks). Below these fields is a "Test Connection" button.

Table 3-52 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 3-52 Synchronization mode

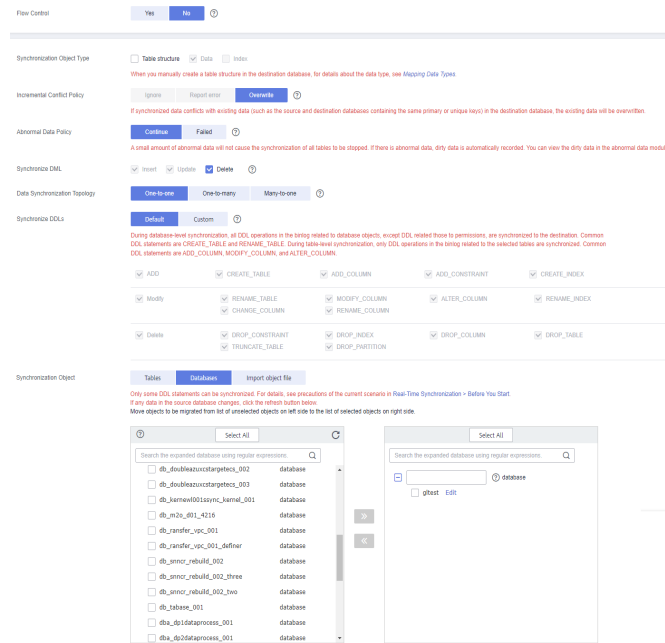
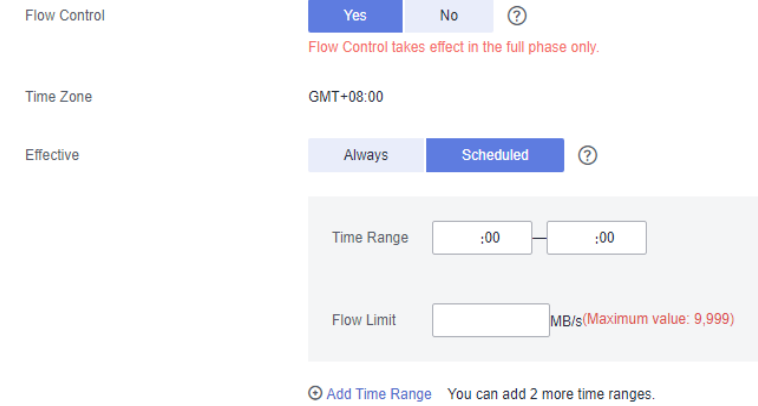



Table 3-53 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of three time ranges can be set, and they cannot overlap. <p>The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.</p> <p>Figure 3-53 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

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

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

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

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

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

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

Figure 3-54 Processing data

| Belonged Database/Table | New name | Column Name | Operation Type | Type | Operation |
|-------------------------|---------------|-------------|----------------|------|-----------|
| db1 (ack_011) | db1 (ack_011) | --- | --- | --- | Add |
| db1 (ack2) | db1 (ack2) | --- | --- | --- | Add |
| db1 (ack3) | db1 (ack3) | --- | --- | --- | Add |

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

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

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

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

Figure 3-55 Pre-check

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

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

NOTE

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

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

Figure 3-56 Task startup settings

Start Time: Start upon task creation Start at a specified time (?)

Send Notifications: (?)

* SMN Topic: (?)

Delay Threshold (s): (?)


Data Exception Notification:

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

Table 3-54 Task startup settings

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

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

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

----End

3.6 From MySQL to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-55 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) Primary/Standby |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-56 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. |

| 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|-----------------|---|
| Source database | <ul style="list-style-type: none"> ● The 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 a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. – If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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 database must be a primary/standby GaussDB(for MySQL) instance.• 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 a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination 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. ● 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. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● If you use additional columns and the number of columns in a single table exceeds 500, adding additional columns may lead to the number of columns in a table to reach an upper limit and cause the task to fail. ● Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |

Procedure

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 3-57 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message in an orange box: "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-8131"), and "Description" (a text area with a character count of "0/256"). A small explanatory text 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."

Table 3-57 Task information

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

- Synchronization instance details

Figure 3-58 Synchronization instance details

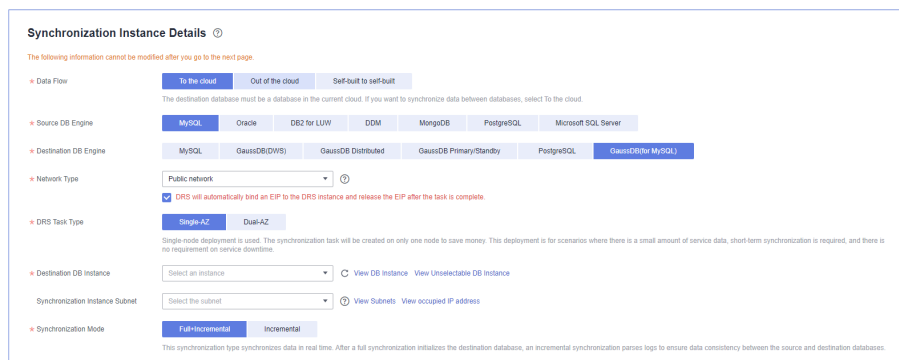


Table 3-58 Synchronization instance settings

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

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

| Parameter | Description |
|----------------------|--|
| Synchronization mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental <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> |


- Task Type

Figure 3-59 Task type



Table 3-59 Task type information

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

| Parameter | Description |
|-----------|--|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-60 AZ</p>  |

- Enterprise Project and Tags

Figure 3-61 Enterprise projects and tags




Table 3-60 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-62 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 Username

Database Password

SSL Connection

● Test successful

Table 3-61 Source database settings

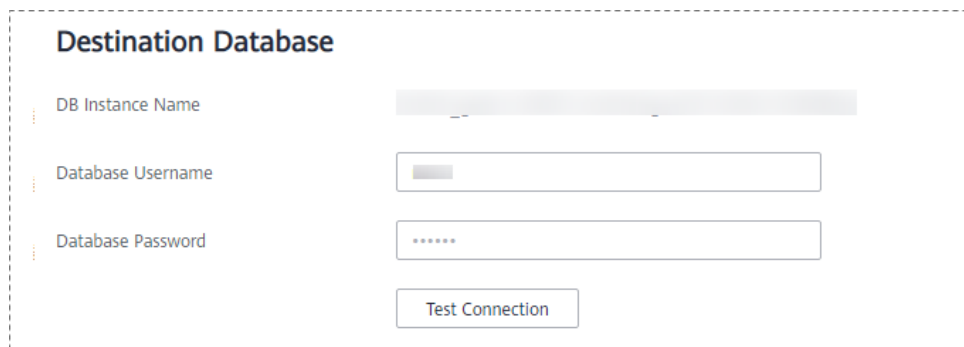
| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source database. |
| Database Password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

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

 **NOTE**

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

Figure 3-63 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 3-62 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 3-64 Synchronization mode

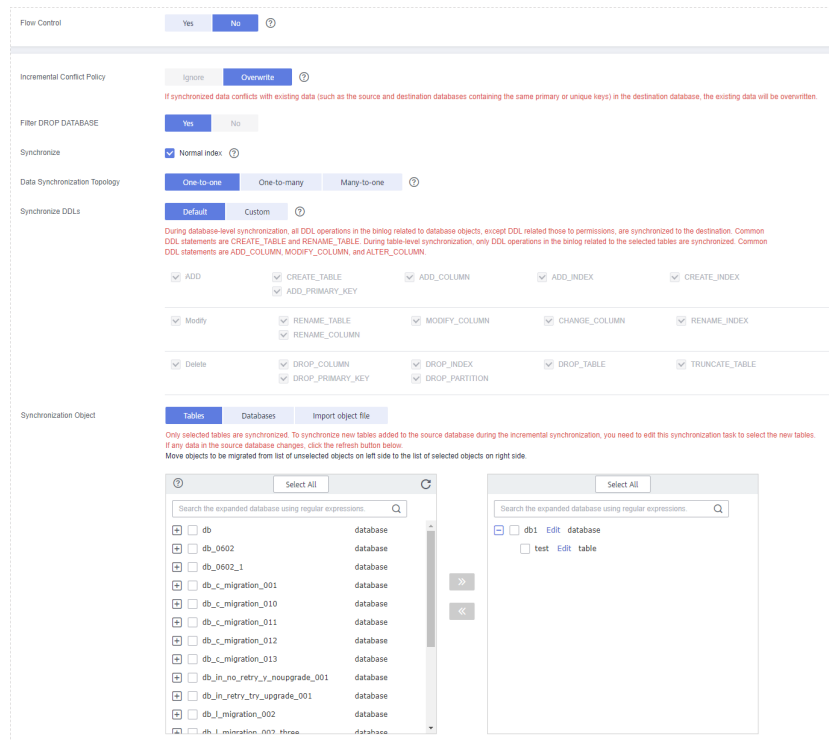
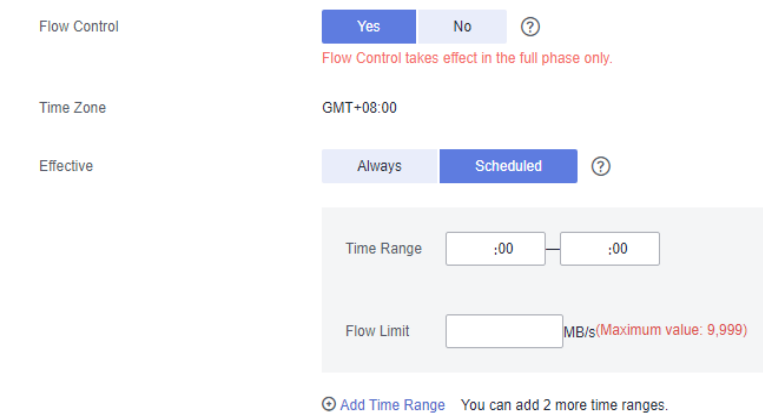



Table 3-63 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-65 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-------------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Whether to synchronize normal indexes.</p> <p>By default, DRS synchronizes the primary key or unique index. A normal index refers to an index other than the primary key or unique index. If you select normal index, all indexes will be synchronized. If you do not select normal index, only the primary key and unique index will be synchronized.</p> |
| Start Point | <p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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>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. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

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

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

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

Figure 3-66 Processing data



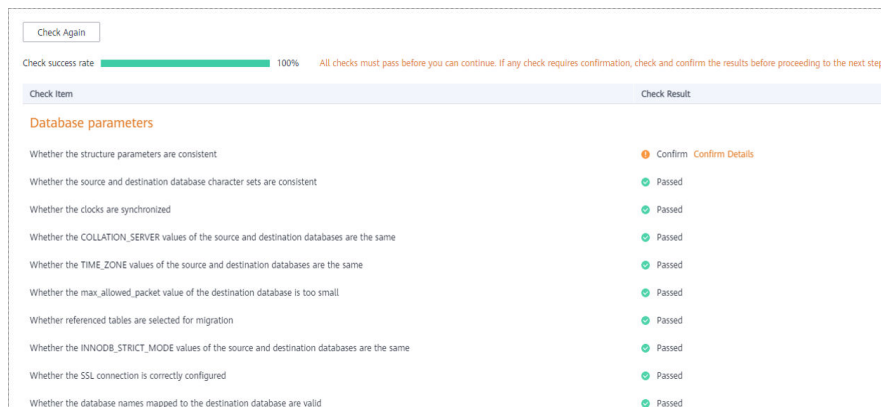
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-67 Pre-check



NOTE

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

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

Figure 3-68 Task startup settings

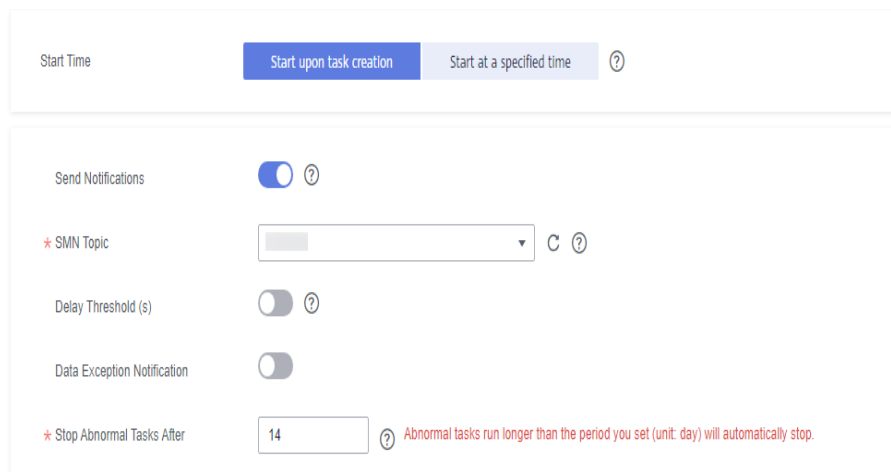



Table 3-64 Task startup settings

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

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

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

----End

3.7 From MySQL to MariaDB

Supported Source and Destination Databases

Table 3-65 Supported databases

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

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-66](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-66 Database account permission

| Type | Full Synchronization | Full+Incremental Synchronization and Incremental Synchronization |
|----------------------|---|---|
| Source database user | The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT | The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT |

| Type | Full Synchronization | Full+Incremental Synchronization and Incremental Synchronization |
|---------------------------|---|--|
| Destination database user | The user must have the following minimum permissions: The root account of RDS for MariaDB has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES | |

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 3-67 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-67 Supported synchronization objects

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

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-68](#).

Table 3-68 Precautions

| 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. • During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. • 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"> ● 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 and MariaDB system databases) 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. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see What Are Syntax Differences Between MySQL or MariaDB Versions? ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● 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 MySQL database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. |

| 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. ● 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. ● DDL operations are not supported during full synchronization. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? |

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

Procedure

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

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

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

- Task information description

Figure 3-69 Synchronization task information

Table 3-69 Task information

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

- Synchronization instance information

Figure 3-70 Synchronization instance information

Table 3-70 Synchronization instance settings

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


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

- AZ

Figure 3-71 AZ



Table 3-71 Task AZ

| Parameter | Description |
|-----------|--|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-72 AZ</p>  |

- Enterprise Project and Tags

Figure 3-73 Enterprise projects and tags




Table 3-72 Enterprise Project and Tags

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

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

 **NOTE**

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

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

- Source database information

Figure 3-74 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 3-73 Source database settings

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

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

 **NOTE**

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

- Destination database information

Figure 3-75 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-74 Destination database settings

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

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

NOTE

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

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

Figure 3-76 Synchronization Mode

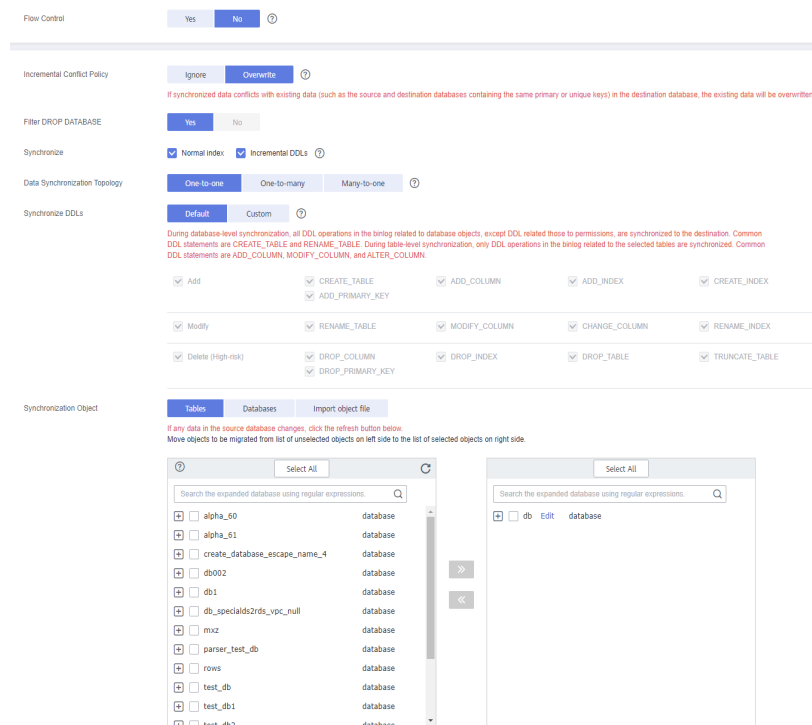
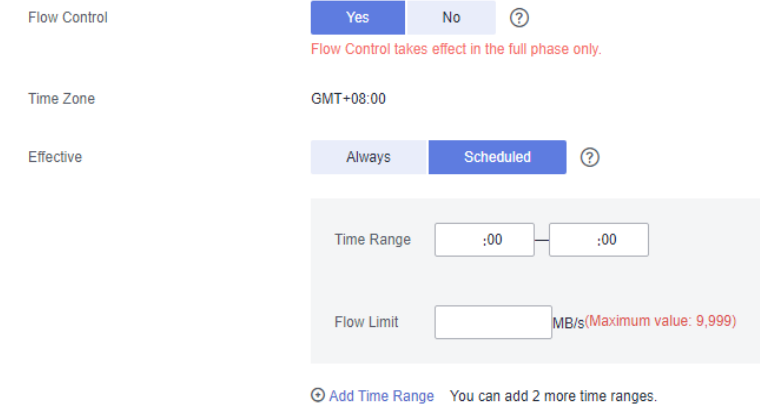



Table 3-75 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-77 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-------------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize 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 position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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> |

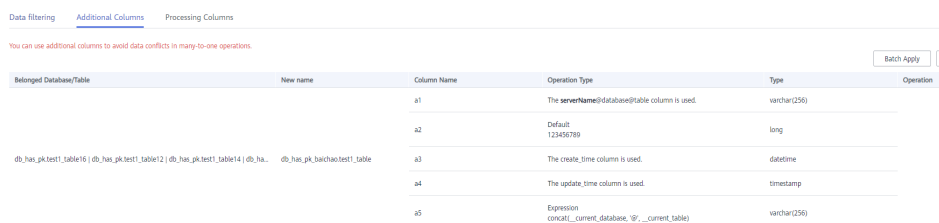
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: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

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

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

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

Figure 3-78 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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

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

 **NOTE**

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

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


Figure 3-79 Task startup settings

Table 3-76 Task startup settings

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

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

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

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

----End

3.8 From PostgreSQL to PostgreSQL

Supported Source and Destination Databases

Table 3-77 Supported databases

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

Supported Synchronization Objects

Table 3-78 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-78 Supported synchronization objects

| Type | Notes |
|---------|---|
| Objects | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. Multiple DRS tasks are required to synchronize multiple databases. ● Supported field types: Digit, currency, character, binary, date/time, boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range. ● Scope of full synchronization <ul style="list-style-type: none"> - The following objects are supported in the database-level synchronization: schemas, tables, indexes, constraints, views, materialized views, sequences, stored procedures, rules, triggers, foreign keys, sorting rules, plug-ins, code conversion information, aggregate functions, operators, statistics extension, conversion information, text search configurations, functions, data types, type conversion, users, event triggers, text search parsers, and text search templates During the table-level synchronization, only tables, views, materialized views, sequences, users, and common indexes can be synchronized. During object file import, tables can be synchronized. - Objects that are not supported: system schemas (schemas starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random), system catalogs, system users, tablespaces, foreign-data wrappers, foreign servers, user mappings, publications, and subscriptions |

| 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. After being synchronized to the destination database, the superuser user becomes a common user and is granted the root permissions by default (supported only when the destination database version is RDS for PostgreSQL 11.11 or later). <ul style="list-style-type: none"> • Scope of incremental synchronization <ul style="list-style-type: none"> - Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. - Some DDL statements can be synchronized, including TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA, CREATE TABLE, DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, ALTER VIEW, COMMENT ON COLUMN, COMMENT ON TABLE, |

| Type | Notes |
|------|---|
| | <p>COMMENT ON SCHEMA, COMMENT ON SEQUENCE, COMMENT ON INDEX, and COMMENT ON VIEW. During table-level synchronization, only the following DDL operations can be synchronized: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, COMMENT ON COLUMN, COMMENT ON TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME).</p> <ul style="list-style-type: none"> - Not synchronized: DML statements of unlogged tables and temporary tables <p>NOTE The source database captures DDL statements using event triggers and records them in specific tables, so you need to create event triggers and functions in the source database in advance. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</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 3-79](#). 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-79 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> <drs_instance_ip>/32 <Authentication mode></code> before all configurations in the pg_hba.conf file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. 2. Run <code>select pg_reload_conf();</code> 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: <ul style="list-style-type: none"> - If the destination database is not PostgreSQL, the CREATEDB permission is required. - If the destination database is PostgreSQL, the CONNECT and CREATE permissions on PostgreSQL databases and the USAGE and CREATE permissions on public schemas are required. ● Table-level: <ul style="list-style-type: none"> - To synchronize databases, the CREATEDB permission is required. - To synchronize a schema, the CONNECT and CREATE permissions for the database that contains the schema are required. - To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the object are required. ● Synchronization user: The CREATEROLE permission is required. ● Synchronization user permissions: The default privilege cannot be modified. Otherwise, the object permissions of the destination database may be inconsistent with those of the source database. <p>NOTE</p> <ul style="list-style-type: none"> ● To synchronize event triggers, text search parsers, and text search templates, the destination database version must be RDS for PostgreSQL 11.11 or later, and the destination database user must be user root or a member of user root. ● If the destination database version is PostgreSQL 15 and no database is created on the destination database for database-level synchronization and table-level synchronization, full or incremental synchronization may fail due to lack of the USAGE and CREATE permissions of the public schema. | |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a

synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 3-80 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 the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database cannot contain objects with the same type and name as the objects to be synchronized, |

| 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. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. |

| Type | Constraints |
|-----------------------------|---|
| | <ul style="list-style-type: none"> - If you choose to synchronize DDL statements, ensure that the DDL statements executed on the source database are compatible with the destination database. <p>NOTE DDL statements are captured using event triggers in the source database, recorded in a specific table, and then synchronized to the destination database. You need to create event triggers and functions in the source database before starting a task. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</p> |
| Full synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. |
| 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. |

| Type | Constraints |
|-----------------|---|
| Stopping a task | <ul style="list-style-type: none"> ● Stop a task normally: <ul style="list-style-type: none"> - The destination database sequence value is automatically reset. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. If users are synchronized, the user memberships are automatically synchronized after the task is complete. - When a full synchronization task is complete, the sequence values are compared based on logical consistency. That is, if an auto-increment sequence value is used only when the sequence value of the destination database is greater than or equal to that of the source database, or an auto-decrement sequence value is used only when the sequence value of the destination database is less than or equal to that of the source database, the sequence values in the source database are consistent with those in the destination database. - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. - If the value of session_replication_role of the destination database is replica when the full+incremental synchronization task is complete, change the value to the original one. ● Forcibly stop a task: <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is <code>drs_unique_ID</code>. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |
| Troubleshooting | <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting. |

Procedure

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

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

Figure 3-80 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 3-81 Task information

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

- Synchronization instance details

Figure 3-81 Synchronization instance details




Table 3-82 Synchronization instance settings

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

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. – Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 3-82 Task type



Table 3-83 Task type information

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

- Enterprise Project and Tags

Figure 3-83 Enterprise projects and tags




Table 3-84 Enterprise Project and Tags

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

NOTE

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

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

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

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

NOTE

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

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

Figure 3-84 Self-build on ECS - source database information

Source Database

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

Database Type: Self-built on ECS RDS DB Instance

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password:

SSL Connection:

Table 3-85 Self-build on ECS - source database information

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

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

 **NOTE**

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

- Scenario 2: RDS DB instance - source database configuration

Figure 3-85 RDS DB instance - source database information

Source Database

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

Database Type: Self-built on ECS RDS DB Instance

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

Database Username:

Database Password:

SSL Connection:

Table 3-86 RDS DB instance - source database information

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

Figure 3-86 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 3-87 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

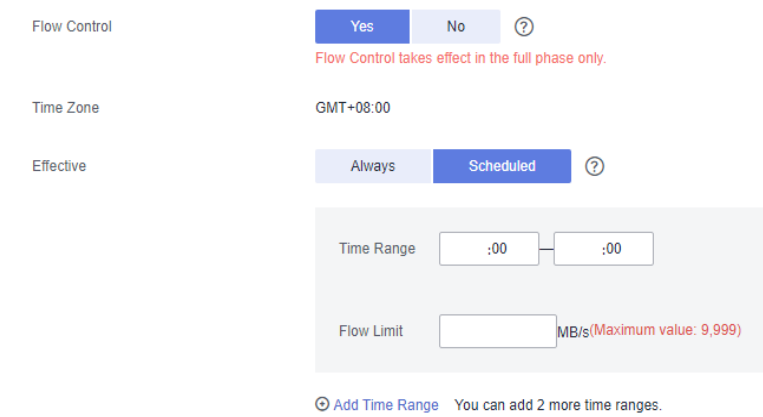
The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

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

Figure 3-87 Synchronization Object

The screenshot displays the "Synchronization Object" configuration page. It includes sections for "Flow Control" (Yes/No), "Incremental Conflict Policy" (Ignore/Report error/Overwrite), "Synchronize" (Normal index/Incremental DDLs/Populate materialized views), "Take Snapshot" (No/Yes), "Concurrent Full Export Threads" (set to 4), "Concurrent Full Import Threads" (set to 8), and "Synchronization Object" (Tables/Databases/Import object file). The "Synchronization Object" section shows a list of databases with checkboxes for selection.

Table 3-88 Synchronization Object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-88 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|--------------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| Synchronize | <p>Available options: 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> |
| Concurrent Full Export Threads | Number of threads for exporting objects and data during full synchronization. |
| Concurrent Full Import Threads | Number of threads for importing objects and data during full synchronization. |


| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • Database-level synchronization: In full synchronization, the selected databases and the inventory data of the database objects are synchronized. In incremental synchronization, the DML and some DDL statements of all tables except unlogged tables and temporary tables are synchronized. • Table-level synchronization: In full synchronization, the inventory data of the selected tables, sequences, views, or materialized views is synchronized. In incremental synchronization, the DML and some DDL statements of the selected tables are synchronized. • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, after the synchronization, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |
| Synchronize Account | <p>During the synchronization, you can synchronize accounts based on your service requirements. For details, see Table 3-89.</p> |

Figure 3-89 Synchronize Account




Table 3-89 Accounts and permissions to be synchronized

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

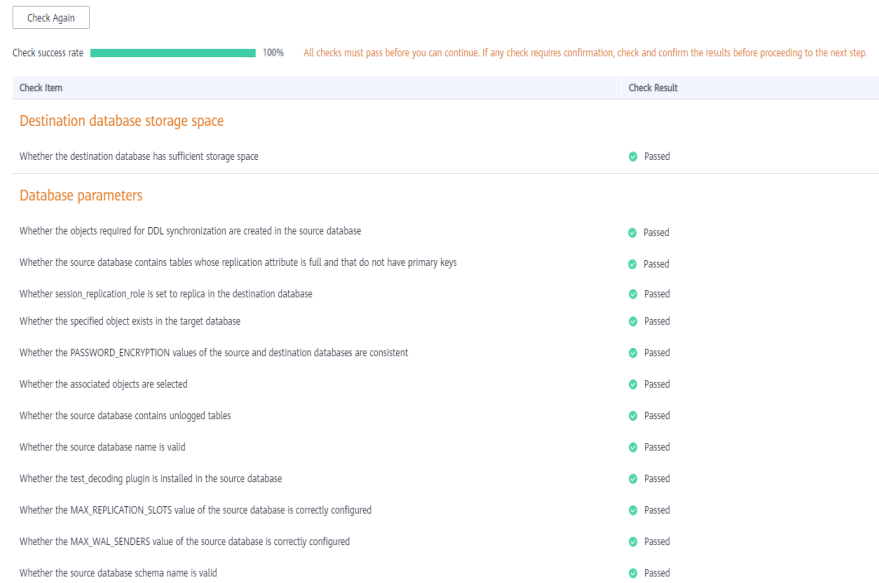
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-90 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 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-91 Task startup settings

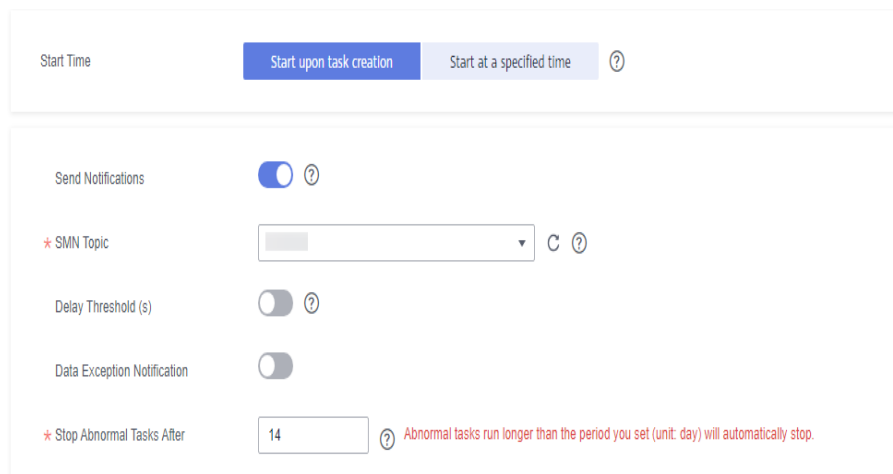



Table 3-90 Task startup settings

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

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

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

----End

3.9 From PostgreSQL to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-91 Supported databases

| Source DB | Destination DB |
|---|--|
| <ul style="list-style-type: none"> • On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, 13 and 14) • ECS 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 3-92](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-92 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> - GaussDB(DWS)-compatible PostgreSQL data types, such as tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz and interval, are supported. - XML, line, and domain synchronization is not supported. ● Scope of full synchronization <ul style="list-style-type: none"> - Schemas, tables, indexes, constraints, sequences, and user-defined types During table-level synchronization, only tables and sequences can be synchronized. - Not supported: system schemas (any schema starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random), and system catalogs. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain +"'%'\<>. The schema name and table name cannot contain ".'\<>. The column name cannot contain "'". The column name cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. ● Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers. ● Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> ● Scope of incremental synchronization <ul style="list-style-type: none"> - Supported: some DML statements, including INSERT, UPDATE, and DELETE - Not supported: DDL statements, DML statements of unlogged tables and temporary tables - DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. |

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 3-93](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

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

Table 3-93 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</code> <code><Authentication mode></code> before all configurations in the <code>pg_hba.conf</code> file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. 2. Run <code>select pg_reload_conf();</code> in the source database as user SUPERUSER, or restart the DB instance to apply the changes. |

| 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. | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 3-94 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The partition table trigger of the source database cannot be set to disable. - For a full synchronization task, the source database can be a standby database, but hot_standby_feedback must be set to on. For an incremental synchronization task, the source database cannot be a standby database. - To perform incremental synchronization: The pg_hba.conf file of the source database contains the following configuration: <div style="background-color: #f0f0f0; padding: 2px; margin: 5px 0;"> <pre>host replication all 0.0.0.0/0 md5</pre> </div> 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"> - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. |
| 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. – The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Procedure

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 3-92 Synchronization task information

Table 3-95 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| 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 3-93 Synchronization instance details

Table 3-96 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 | <p>Available options: VPC, Public network and VPN or Direct Connect. VPC is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | The destination is a GaussDB(DWS) instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. |

- AZ

Figure 3-94 AZ




Table 3-97 Task AZ

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

- Enterprise Project and Tags

Figure 3-95 Enterprise projects and tags

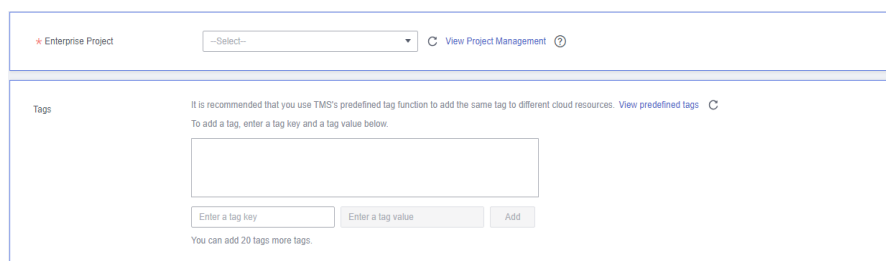


Table 3-98 Enterprise Project and Tags

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

 **NOTE**

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

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

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

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

 **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-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 3-99 Self-built on ECS - source database information

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

Test Connection: Test successful

Table 3-100 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 3-98 Destination database information

Destination Database

DB Instance Name

Database Username


Database Password 

Table 3-101 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 3-99 Synchronization Mode

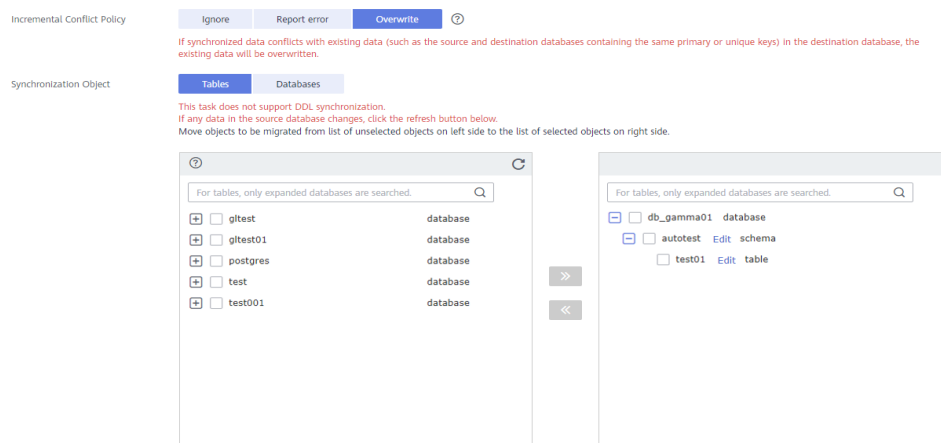



Table 3-102 Synchronization object

| 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. If you select Ignore, data in the source database may be inconsistent with that in the destination database. Report error The synchronization task will be stopped and fail. Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |

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

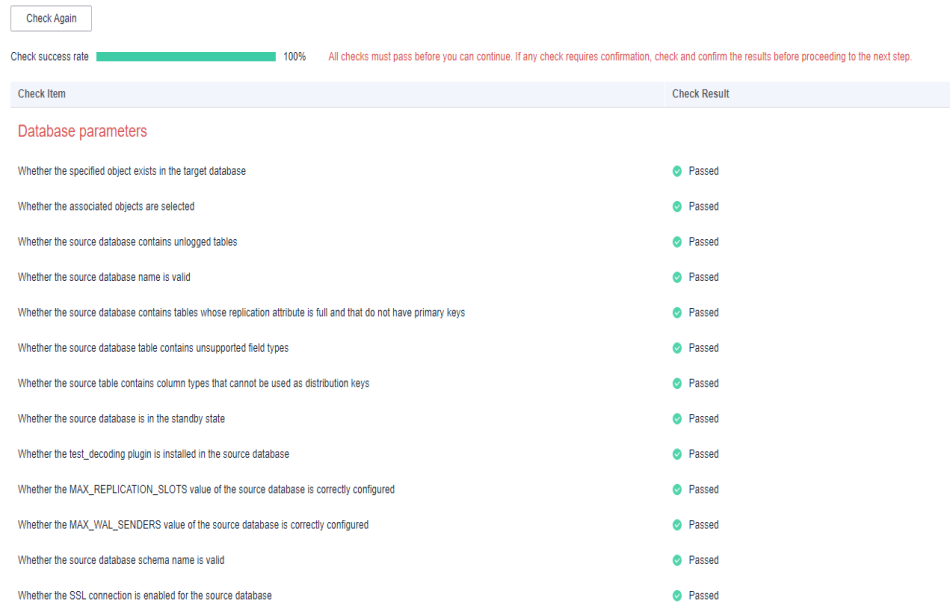
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-101 Task startup settings

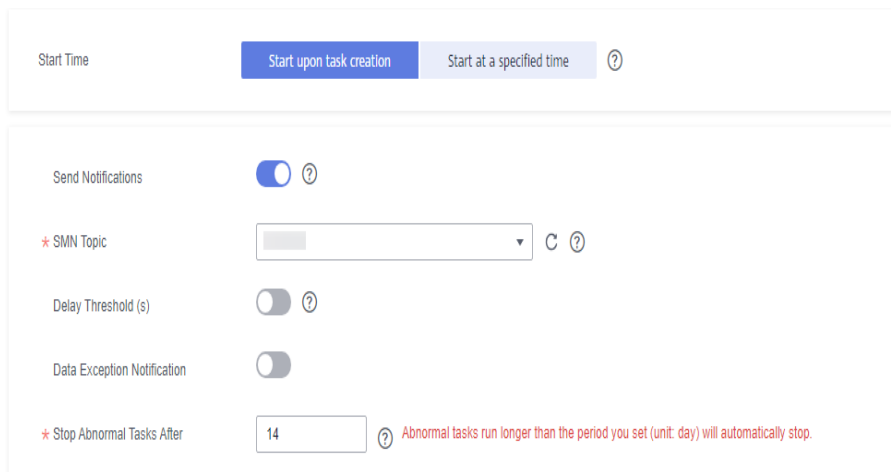



Table 3-103 Task startup settings

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

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

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

----End

3.10 From PostgreSQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-104 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 3-105](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-105 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> - GaussDB-compatible PostgreSQL data types, such as tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestamptz and interval, are supported. - XML, line, domain, and self-built data type synchronization is not supported. ● Scope of full synchronization <ul style="list-style-type: none"> - Supported: schemas, tables, primary 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"). - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain "+" %?<>, the schema name and table name cannot contain ".!<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). ● Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. ● Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> ● Scope of incremental synchronization <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- DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. |

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 3-106](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

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

Table 3-106 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, and the SELECT permission for sequences | <p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the pg_hba.conf file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes. |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. If <code>gs_loader</code> is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | |

Suggestions

⚠ CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 3-107 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 dependencies: tables referenced by views, views referenced by views, tables referenced by primary and foreign keys, parent tables referenced by child tables, partitioned tables referenced by sub-partitioned tables, and sequences referenced by auto-increment columns.</p> ● Destination database parameter requirements: <ul style="list-style-type: none"> - The max_prepared_transactions value of the destination database must be greater than that of the source database. - The max_worker_processes value of the destination database must be greater than that of the source database. - The max_locks_per_transaction value of the destination database must be greater than that of the source database. |

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

| Type | Constraints |
|----------------------|---|
| | <p>name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key.</p> <ul style="list-style-type: none"> - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. |
| 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 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"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. ● Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Procedure

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

Figure 3-102 Synchronization task information

Table 3-108 Task information

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

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

- Synchronization instance details

Figure 3-103 Synchronization instance details

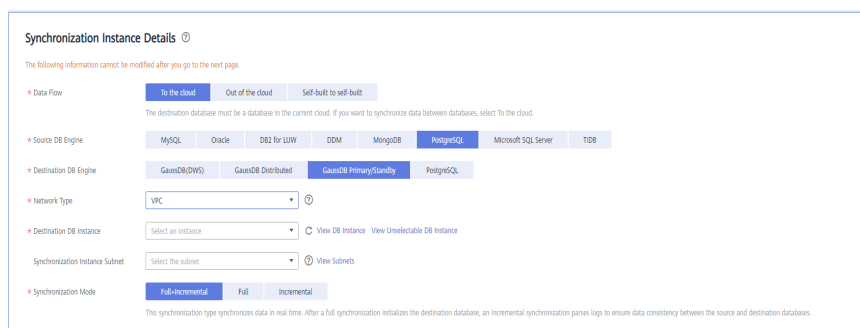


Table 3-109 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|--|
| Destination DB Instance | An available GaussDB primary/standby instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full This synchronization mode is suitable for scenarios where service interruption is acceptable. Database objects and data, including tables, views, and stored procedures, from non-system databases can be synchronized to the destination all at once. - Incremental In this mode, incremental data generated on the source database is continuously synchronized to the destination database through log parsing. |

- AZ

Figure 3-104 AZ




Table 3-110 Task AZ

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

- Enterprise Project and Tags

Figure 3-105 Enterprise projects and tags




Table 3-111 Enterprise Project and Tags

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

NOTE

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

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

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

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

 **NOTE**

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

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

Figure 3-106 Self-build on ECS - source database information

Source Database

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

Database Type: Self-built on ECS RDS DB instance

VPC: View VPC

Subnet: View Subnets

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password:

SSL Connection:

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

| Parameter | Description |
|---------------|---|
| Database Type | Select Self-built on ECS . |
| VPC | A dedicated virtual network in which the source database is located. It isolates networks for different services. You can select an existing VPC or create a VPC. |

| Parameter | Description |
|---------------------------|---|
| 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 3-107 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, there are two tabs: 'Self-built on ECS' and 'RDS DB Instance'. The 'RDS DB Instance' tab is selected. Underneath, there is a dropdown menu for 'DB Instance Name' with a search icon and two links: 'View DB Instance' and 'View Unselectable DB Instance'. Below the dropdown are input fields for 'Database Username' and 'Database Password'. At the bottom, there is a 'Test Connection' button.

Table 3-113 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. |

| Parameter | Description |
|-------------------|---|
| Database Password | The password for the database username. |

Figure 3-108 Destination database information

The screenshot shows a form titled "Destination Database" with the following fields and a button:

- 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 3-114 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 3-109 Synchronization mode

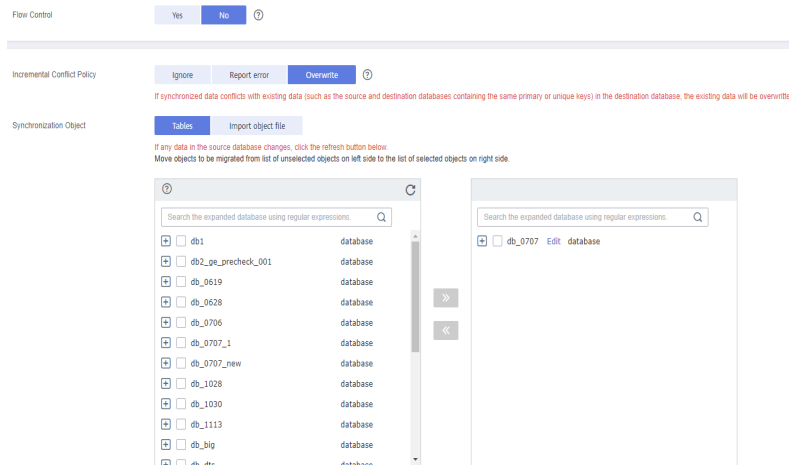
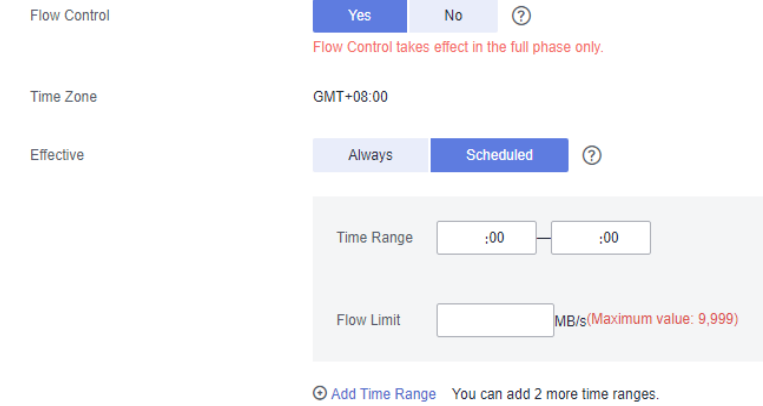



Table 3-115 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-110 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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. ● When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

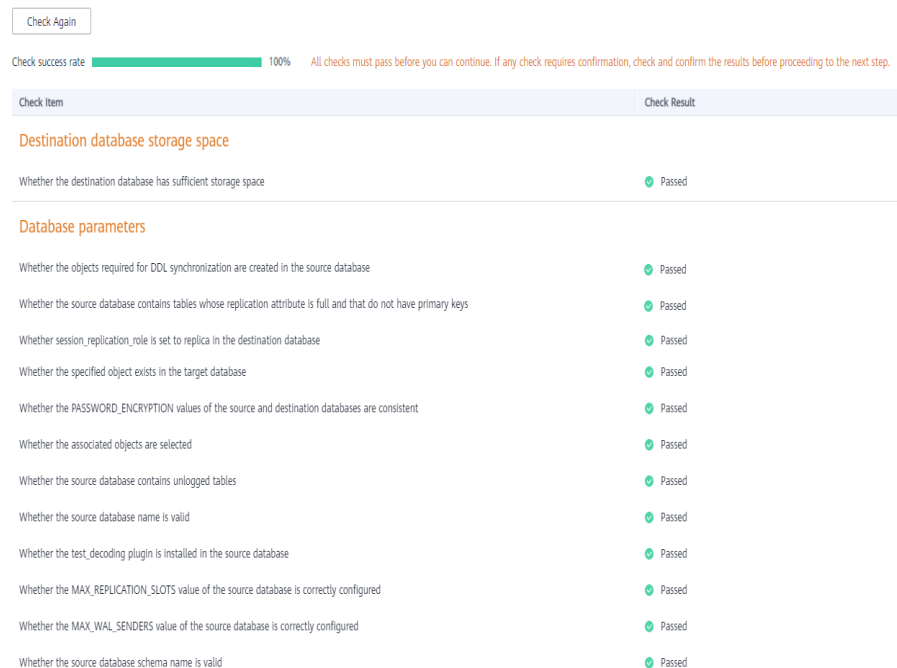
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-111 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-112 Task startup settings





Table 3-116 Task startup settings

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

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

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

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

----End

3.11 From PostgreSQL to GaussDB Distributed

Supported Source and Destination Databases

Table 3-117 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 3-118 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-118 Supported synchronization objects

| Type | Precautions |
|---------|--|
| Objects | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple DRS tasks. ● Supported field types: <ul style="list-style-type: none"> – GaussDB-compatible PostgreSQL data types, such as tinyint, smallint, int, bigint, numeric, decimal, char, bpchar, varchar, text, date, time, timetz, timestamp, timestampz and interval, are supported. – XML, line, 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"). – A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain "+" %?\<>, the schema name and table name cannot contain ".\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). ● Table: Temporary tables are not synchronized. Only table-level synchronization is supported. Constraints on tables can be synchronized together. Common indexes, rules, and triggers of tables cannot be synchronized. Only objects in the primary table can be synchronized. All data in the partition table will be written to the primary table. ● Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized. <ul style="list-style-type: none"> ● Scope of incremental synchronization |

| Type | Precautions |
|------|---|
| | <ul style="list-style-type: none">- Supported: some DML statements, including INSERT, UPDATE, and DELETE- Not supported: DDL statements, DML statements of unlogged tables and temporary tables- DML synchronization of partitioned tables: If a partition is a new one, the partition must be newly created and the partition name must be unique. This partition can be deleted from the source database only after all DML data is synchronized to the destination database. |

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 3-119](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

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

Table 3-119 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, and the SELECT permission for sequences | <p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE</p> <ul style="list-style-type: none"> • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: <ol style="list-style-type: none"> 1. Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32</code> <code><Authentication mode></code> before all configurations in the pg_hba.conf file of the source database. For details about the authentication mode, see pg_hba.conf in the official document of PostgreSQL. <code>scram-sha-256</code> is a common authentication mode. 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes. |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If <code>gs_loader</code> is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 3-120 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 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: <code>i_+hash value +original index name (which may be truncated)+_key</code> The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to <code>c_</code> |

| Type | Constraints |
|----------------------|--|
| | <ul style="list-style-type: none"> + hash value + original constraint name (which may be truncated) + _key. - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. If a long transaction is started in the source database, the creation of the logical replication slot will be blocked. As a result, the task fails. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - In a full synchronization for the table structure, the length of the character and character varying types in the source database automatically increases by byte in the destination database (because the length of the destination database is in the unit of byte). By default, the length is increased by four times, and the maximum length is 10485760. - If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails. |
| Full synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. |

| Type | Constraints |
|-----------------------------|---|
| Incremental synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail. • Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • 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"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Procedure

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

Figure 3-113 Synchronization task information

The screenshot shows a web form for creating a synchronization instance. 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 3-121 Task information

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

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

- Synchronization instance details

Figure 3-114 Synchronization instance details

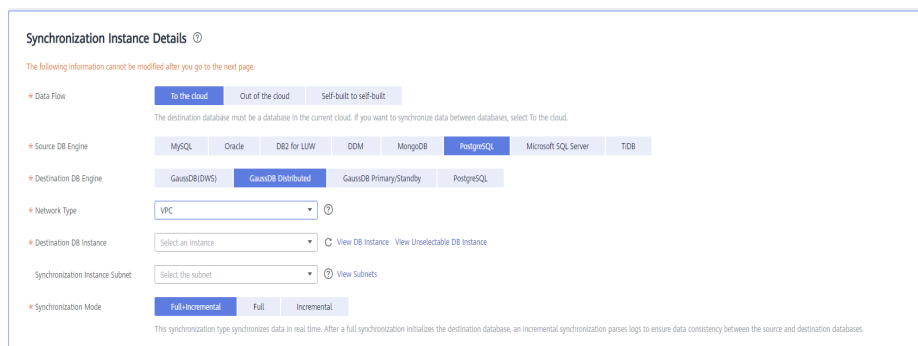


Table 3-122 Synchronization instance information

| Parameter | Description |
|-----------------------|-------------------------------------|
| Data Flow | Select To the cloud . |
| Source DB Engine | Select PostgreSQL . |
| Destination DB Engine | Select GaussDB Distributed . |

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

- AZ

Figure 3-115 AZ



Table 3-123 Task AZ

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

- Enterprise Project and Tags

Figure 3-116 Enterprise projects and tags

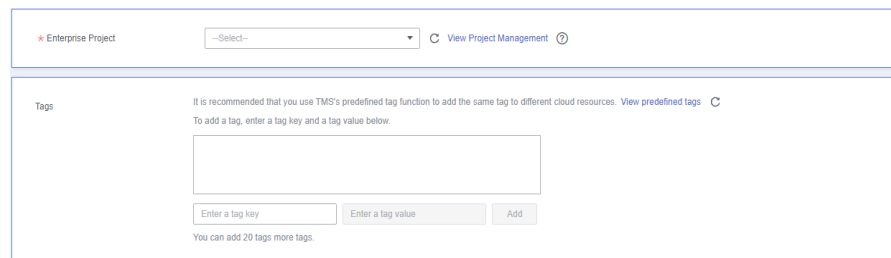


Table 3-124 Enterprise Project and Tags

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

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

 **NOTE**

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

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

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

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

 **NOTE**

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

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

Figure 3-117 Self-build on ECS - source database information

Source Database

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

Database Type: Self-built on ECS RDS DB Instance

VPC: [View VPC](#)

Subnet: [View Subnets](#)

IP Address or Domain Name:

Port:

Database Name:

Database Username:

Database Password: [Show/Hide](#)

SSL Connection:

Table 3-125 Self-build on ECS - source database information

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

NOTE

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

- Scenario 2: RDS DB instance - source database configuration

Figure 3-118 RDS DB instance - source database information

The screenshot shows a configuration form titled "Source Database". At the top, there is a note: "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, there are two radio buttons for "Source Database Type": "Self-built on ECS" and "RDS DB Instance", with the latter selected. Underneath, there is a dropdown menu for "DB Instance Name" with a "View DB Instance" link and a "View Unselectable DB Instance" link. Below that are input fields for "Database Username" and "Database Password" (masked with asterisks). A "Test Connection" button is located at the bottom of the form.

Table 3-126 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 3-119 Destination database information

The screenshot shows a configuration form titled "Destination Database". It contains a dropdown menu for "DB Instance Name", followed by input fields for "Database Username" and "Database Password" (masked with asterisks). A "Test Connection" button is positioned at the bottom of the form.

Table 3-127 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 3-120 Synchronization mode

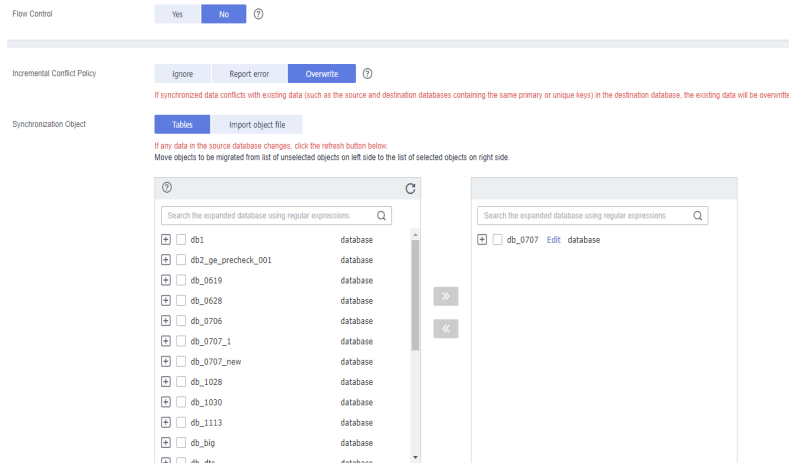
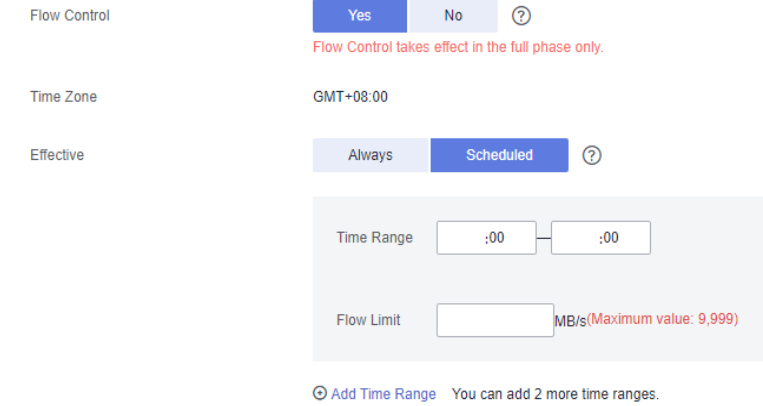



Table 3-128 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-121 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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. ● When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. ● For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

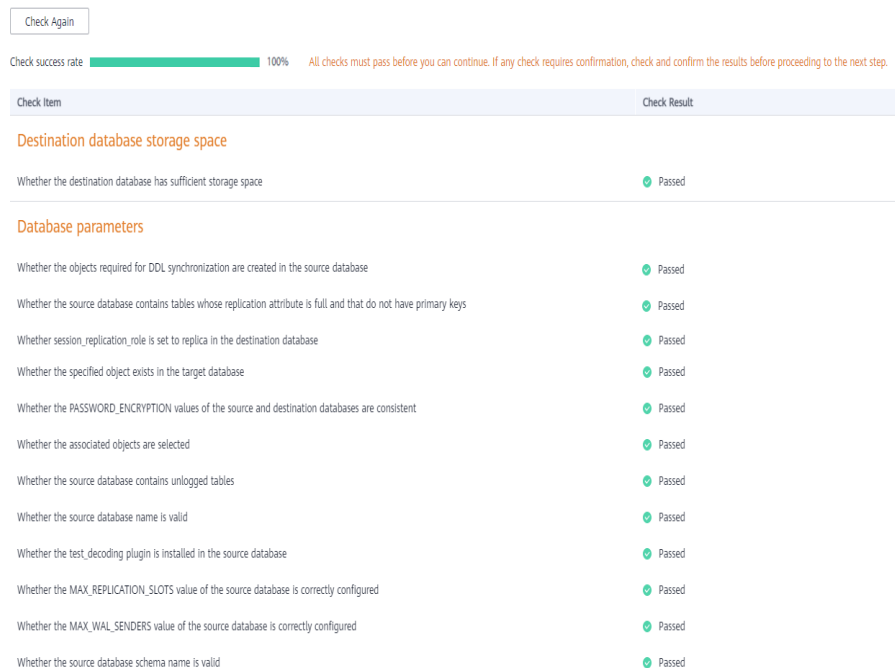
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-122 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 6** On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 3-123 Task startup settings





Table 3-129 Task startup settings

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

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

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

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

----End

3.12 From Oracle to MySQL

Supported Source and Destination Databases

Table 3-130 Supported databases

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

Supported Synchronization Objects

[Table 3-131](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-131 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> • Object level: table level, object file import, • 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. - In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. - Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by primary or foreign keys - Partitions in the table structure cannot be synchronized. Partitioned tables are changed to non-partitioned tables after being synchronized to the destination database. - Tables whose default values contain expression functions cannot be synchronized. - Temporary tables in the source database cannot be synchronized. - An empty source database cannot be synchronized. - Tables with virtual columns in the source database cannot be synchronized. - 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 | Precautions |
|------|--|
| | <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. <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 3-132](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

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

Table 3-132 Database account permission

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

| Type | Full | Incremental and Full+Incremental |
|---------------------------|---|--|
| | | <p>drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p> <ul style="list-style-type: none"> • During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. • Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs. |
| 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.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 3-133 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 precision of the float(p) field in the Oracle database cannot exceed the precision range p: [1, 32]. The float range of MySQL is smaller than that of Oracle. - The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record. - The total index length of columns in the source database cannot exceed the length limit in the destination database. 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: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Restrictions |
|------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <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 data types are incompatible, the synchronization may fail. - The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. - If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code. - After the Oracle table structure is synchronized to the MySQL database, the character set of the table is UTF8MB4. - If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. - If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. - In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. - For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can |

| Type | Restrictions |
|------|--|
| | <p>provide better fault tolerance, load capability, and synchronization experience.</p> <ul style="list-style-type: none"> - If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. - There are some syntax differences between Oracle and MySQL, so the syntax including but not limited to functions, expressions, and referenced system tables, may not be completely converted during the structure synchronization. Therefore, during the synchronization, the structure may exist in the Oracle database but does not exist in the MySQL database, or the syntax exists in the MySQL database but is not converted. As a result, the structure fails to be synchronized. If this happens, you can manually create a table structure in the destination database. - You can enable strong consistency during synchronization to ensure that the commit sequence and atomicity of transactions synchronized to the destination database are the same as those of the source database. However, the performance in this mode is much lower than that in the default mode. - When strong consistency is enabled, out-of-memory (OOM) may occur if the size of committed transactions is greater than 256 MB. - In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. - You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. |

| Type | Restrictions |
|----------------------|---|
| 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 a full synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During a full synchronization, do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● During a full synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. ● During a full synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. ● During a full synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● 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. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● During full data synchronization, a lot of binlogs are generated in the destination database, occupying too much storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. |

| Type | Restrictions |
|-----------------------------|---|
| Incremental synchronization | <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, do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● During an incremental synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. ● During an incremental synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. ● During an incremental synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● During an incremental 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 an incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - 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. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • 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. 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 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 3-124 Synchronization task information

Table 3-134 Task information

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

- Synchronization instance details

Figure 3-125 Synchronization instance details

Table 3-135 Synchronization instance information

| Parameter | Description |
|-----------|------------------------------|
| Data Flow | Select To the cloud . |

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

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 3-126 Task type



Table 3-136 Task type information

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

- Enterprise Project and Tags

Figure 3-127 Enterprise projects and tags

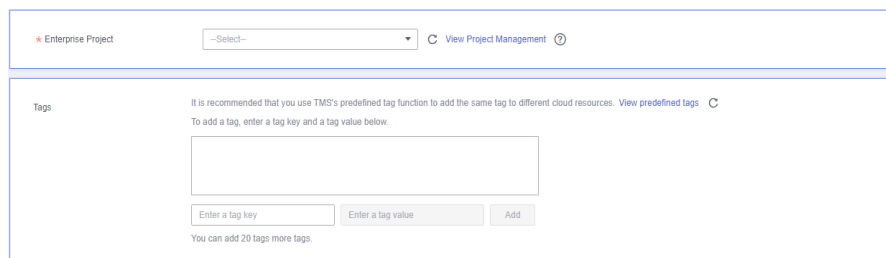


Table 3-137 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-128 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 3-138 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 disabled, your data may be at risk. |

 **NOTE**


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

Figure 3-129 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-139 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 disabled, your data may be at risk. |

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

Figure 3-130 Synchronization mode

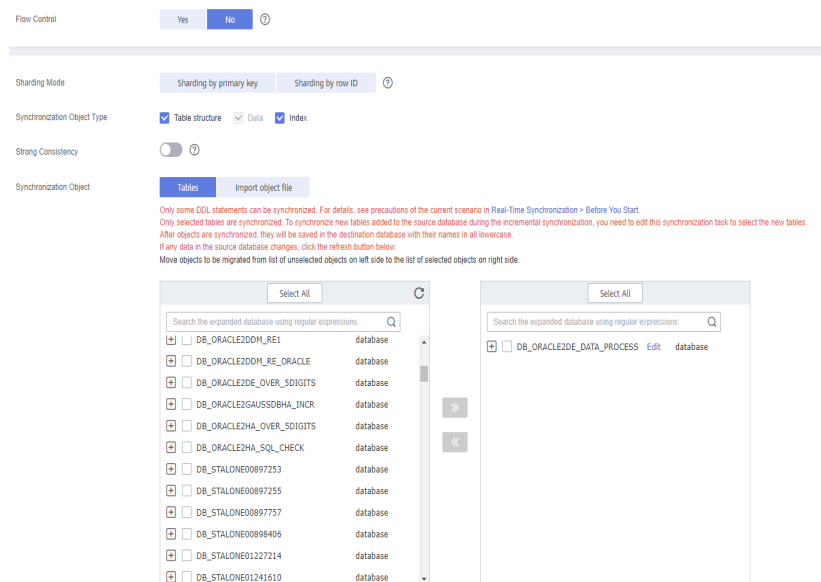
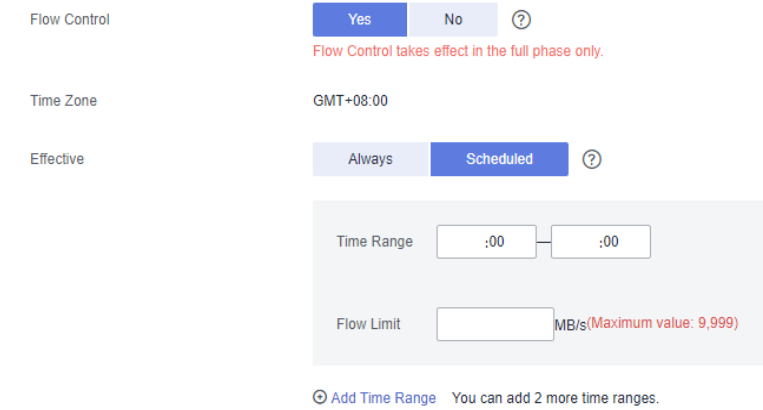



Table 3-140 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-131 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

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

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

Step 5 On the **Process Data** page, 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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

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

 **NOTE**

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

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

Figure 3-132 Task startup settings

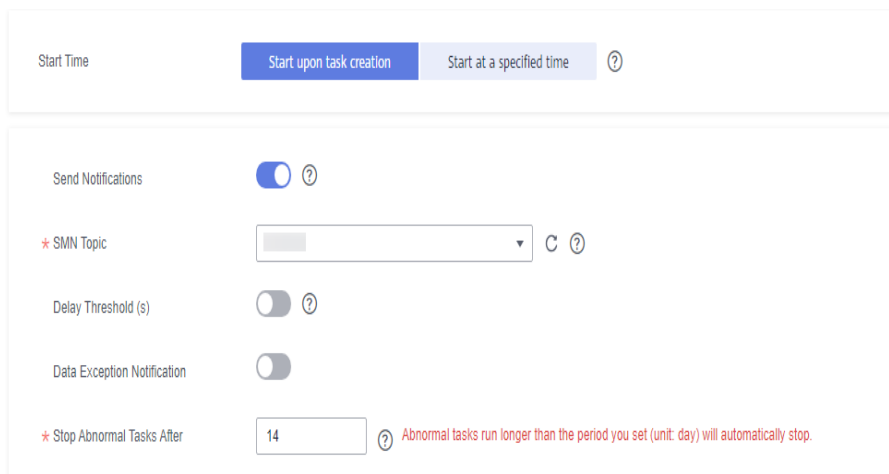



Table 3-141 Task startup settings

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

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

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

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

----End

3.13 From Oracle to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-142 Supported databases

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

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-143 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 ON <userName.tbName> to drsUser). - Full+Incremental synchronization: <ul style="list-style-type: none"> Oracle 12c or later in tenant mode: <p>To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <p>To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB.</p> <p>Oracle 12c or later in non-tenant mode: <p>You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.</p> <p>To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.</p> </p> - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"><li data-bbox="592 412 1342 539">• The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, INDEX, EVENT, RELOAD, CREATE VIEW, CREATE ROUTINE, and TRIGGER. |

| Type | Restrictions |
|------------------------|--|
| 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● 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 expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. |

| 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> |
| <p>Destination database</p> | <ul style="list-style-type: none"> ● The destination database must be a primary/standby GaussDB(for MySQL) instance. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code. ● After the Oracle table structure is synchronized to the MySQL database, the character set of the table is UTF8MB4. ● If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the |

| Type | Restrictions |
|------|---|
| | <p>first time. Otherwise, an error occurs during incremental synchronization.</p> <ul style="list-style-type: none"> ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● 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. ● 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, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● During 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. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If a task is abnormal due to DDL synchronization in other cases, you need to manually execute the DDL operations in the destination database. <ul style="list-style-type: none"> – Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. – 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. – 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 |

| Type | Restrictions |
|------|--|
| | <p>synchronized to the destination database. You are advised to add tables during off-peak hours.</p> <ul style="list-style-type: none"> You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. |

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 3-133 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, there are four main fields:

- Region:** A dropdown menu with a small 'Q' icon to its left.
- Project:** A dropdown menu with a small 'Q' icon to its left.
- * Task Name:** A text input field containing "DRS-6131" and a help icon to its right.
- Description:** A larger text area with a help icon to its right.

 A small "0/256" character count is visible at the bottom right of the description field.

Table 3-144 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 3-134 Synchronization instance details

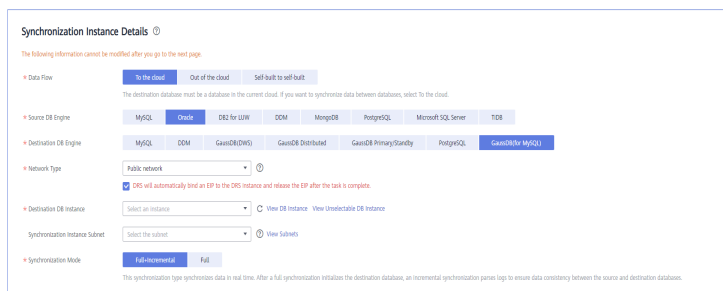


Table 3-145 Synchronization instance settings

| 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 <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | The GaussDB(for MySQL) instance you created. |

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE</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"> 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. |

- Task type

Figure 3-135 Task type



Table 3-146 Task type information

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

- Enterprise Project and Tags

Figure 3-136 Enterprise projects and tags




Table 3-147 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-137 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 | <p>The IP address or domain name of the source database.</p> <p>NOTE For a RAC cluster, use a Scan IP address to improve access performance.</p> |

| Parameter | Description |
|-----------------------|---|
| Port | The port of the source database. Range: 1 – 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| PDB Name | Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **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 3-138 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-149 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 3-139 Synchronization mode

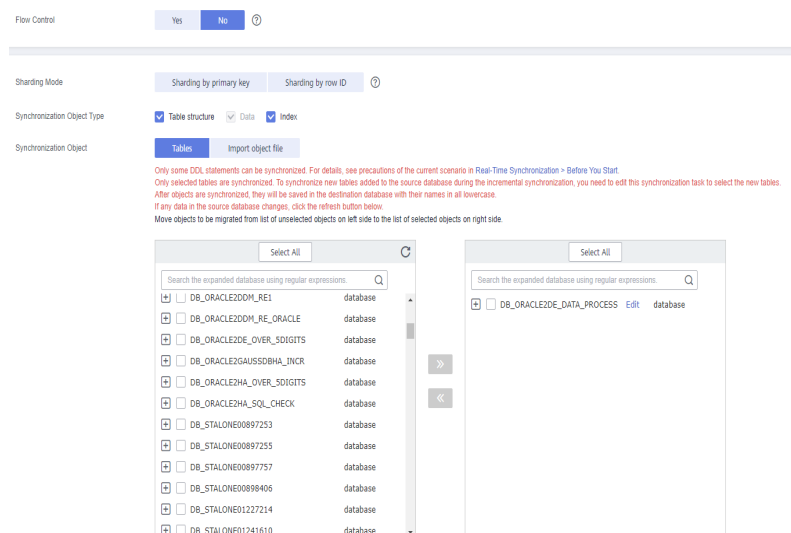
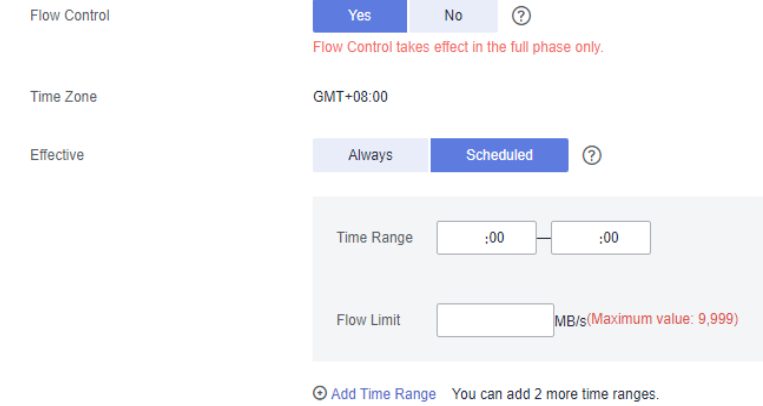



Table 3-150 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 |
|-----------------------------|---|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> ● Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. ● Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations. |
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |

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

Step 5 On the **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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

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

 **NOTE**

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

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

Figure 3-141 Task startup settings





Table 3-151 Task startup settings

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

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

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

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

----End

3.14 From Oracle to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-152 Supported databases

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

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-153 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 ON <userName.tbName> to drsUser). - Full+incremental synchronization and incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"> • The destination database must have the following permissions: <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> - Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> - Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, and function indexes. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute and does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● For the TIMESTAMP WITH TIME ZONE data type, the 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 |

| Type | Restrictions |
|-----------------|--|
| | <p>3. Otherwise, incremental synchronization may fail because all columns cannot be matched.</p> <ul style="list-style-type: none"> • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • Tables with virtual columns in the source database cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. • When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. |
| 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 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 due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. ● 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. |

| Type | Restrictions |
|------|--|
| | <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. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● 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, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● Do not 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 |

| Type | Restrictions |
|------|--|
| | <p>synchronized, so you must manually create the table in the destination database.</p> <ul style="list-style-type: none"> ● During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. ● During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. ● In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● 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, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases |

| Type | Restrictions |
|------|--|
| | <p>requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database.</p> <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. • During incremental synchronization, 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 hidden columns (UNUSED, INVISIBLE). • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. • You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. • When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may occur after the switchover due to DDL changes during replication delay or long transactions. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database |

| Type | Restrictions |
|------|--|
| | sequence value minus the security margin. The default security margin is 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 3-142 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]

Table 3-154 Task information

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

- Synchronization instance details

Figure 3-143 Synchronization instance details

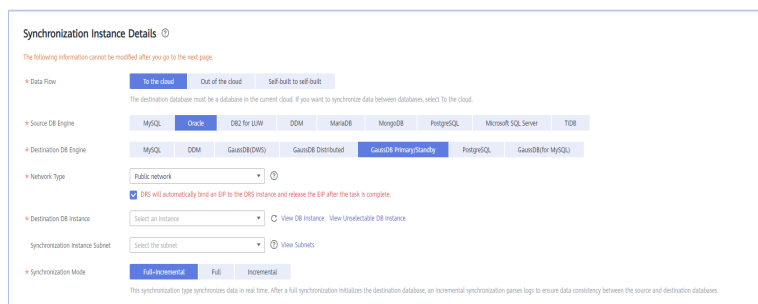


Table 3-155 Synchronization instance settings

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

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . |

- Task type

Figure 3-144 Task type



Table 3-156 Task type information

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

- Enterprise Project and Tags

Figure 3-145 Enterprise projects and tags




Table 3-157 Enterprise Project and Tags

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

 **NOTE**

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

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

Figure 3-146 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 3-158 Source database settings

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

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

Figure 3-147 Destination database information

The screenshot shows a form titled "Destination Database" with three input fields and a button. The fields are labeled "DB Instance Name", "Database Username", and "Database Password". The "Database Password" field is masked with asterisks. Below the fields is a button labeled "Test Connection".

Table 3-159 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 3-148 Synchronization mode

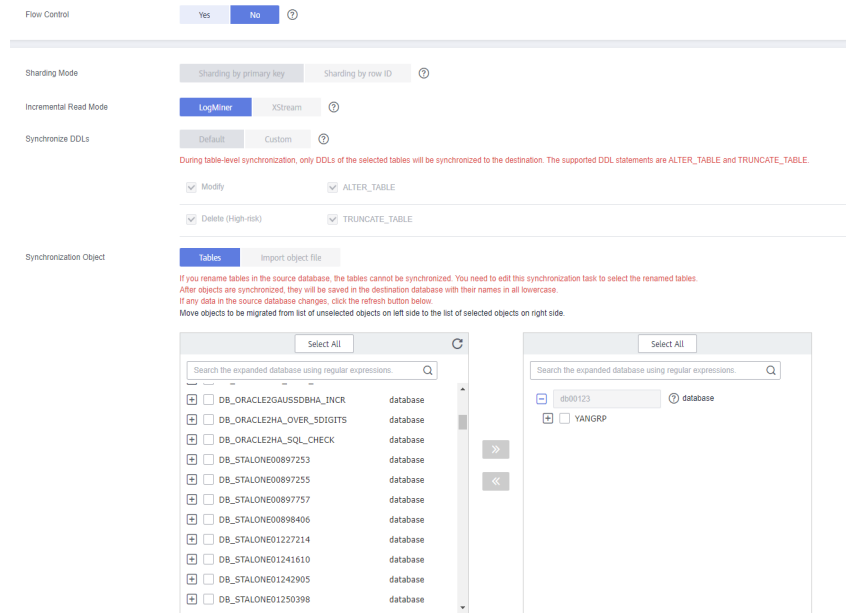
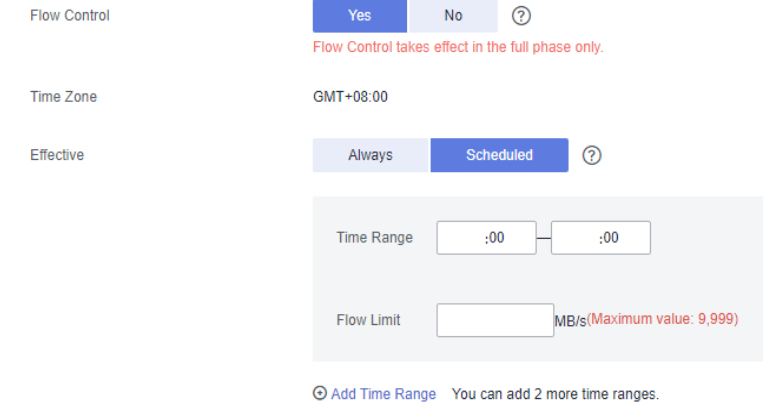



Table 3-160 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-149 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

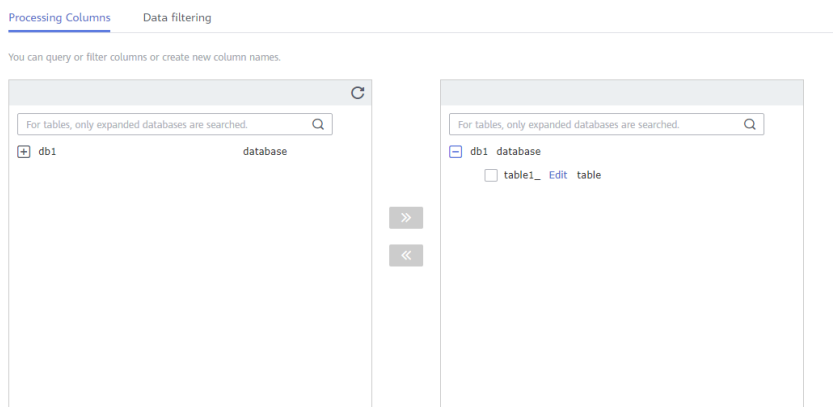
| Parameter | Description |
|-----------------------|--|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations. |
| Incremental Read Mode | <p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. XStream: The XStream interface must be enabled. <p>LogMiner is recommended.</p> |
| Synchronize DDLs | <p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. |

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

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

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

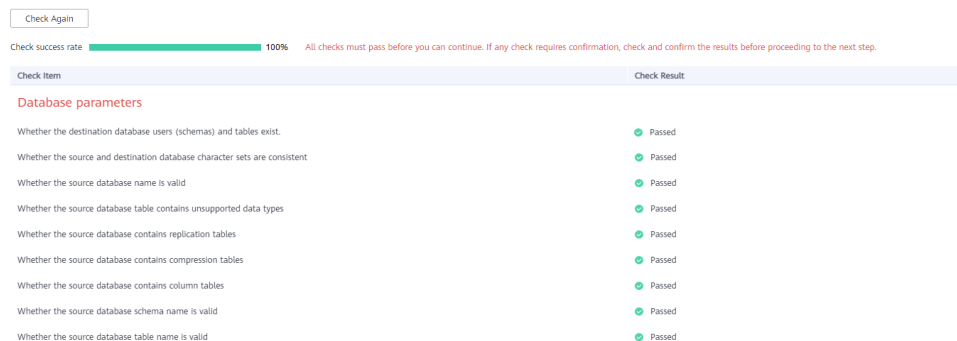
Figure 3-150 Processing data



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

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 3-151 Pre-check



NOTE

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

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

Figure 3-152 Task startup settings





Table 3-161 Task startup settings

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

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

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

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

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

3.15 From Oracle to GaussDB Distributed

Supported Source and Destination Databases

Table 3-162 Supported databases

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

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-163 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 ON <userName.tbName> to drsUser). - Full+incremental synchronization and incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"> • The destination database must have the following permissions: <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> - Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> - Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● 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 full synchronization, the table structure does 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. |

| Type | Restrictions |
|-----------------|--|
| | <ul style="list-style-type: none"> ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. ● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. |
| 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 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 due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. ● 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. |

| Type | Restrictions |
|------|--|
| | <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. • If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. • If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. • 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, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. • Do not 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 |

| Type | Restrictions |
|------|--|
| | <p>synchronized, so you must manually create the table in the destination database.</p> <ul style="list-style-type: none"> ● In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. ● During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. ● During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● 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, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases |

| Type | Restrictions |
|------|--|
| | <p>requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database.</p> <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. • During incremental synchronization, 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 hidden columns (UNUSED, INVISIBLE). • 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. • You are advised to disable the global secondary index (GSI) function for the destination database. Otherwise, incremental synchronization may fail. • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. • You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. • When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may |

| Type | Restrictions |
|------|--|
| | <p>occur after the switchover due to DDL changes during replication delay or long transactions.</p> <ul style="list-style-type: none"> If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. |

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 3-153 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 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, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "Task Name" (a text input field containing "DRS-6131" with a help icon), and "Description" (a text area with a help icon). A small "0/256" character count is visible at the bottom right of the description field.

Table 3-164 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 3-154 Synchronization instance details

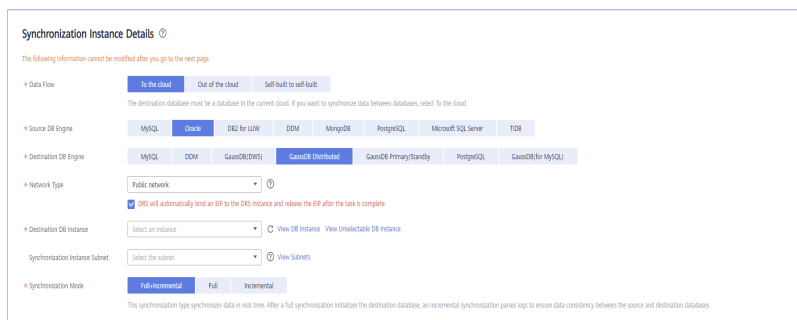


Table 3-165 Synchronization instance settings

| Parameter | Description |
|-------------------------|---|
| Data Flow | Select To the cloud . |
| Source DB Engine | Select Oracle . |
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | The GaussDB distributed instance you created. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Synchronization Mode | The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . |

- Task type

Figure 3-155 Task type




Table 3-166 Task type information

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

- Enterprise Project and Tags

Figure 3-156 Enterprise projects and tags




Table 3-167 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-157 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 3-168 Source database settings

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

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

Figure 3-158 Destination database information

The screenshot shows a form titled "Destination Database" with three input fields and a button. The fields are labeled "DB Instance Name", "Database Username", and "Database Password". The "DB Instance Name" field is a dropdown menu. The "Database Username" and "Database Password" fields are text boxes. Below the fields is a "Test Connection" button.

Table 3-169 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 3-159 Synchronization mode

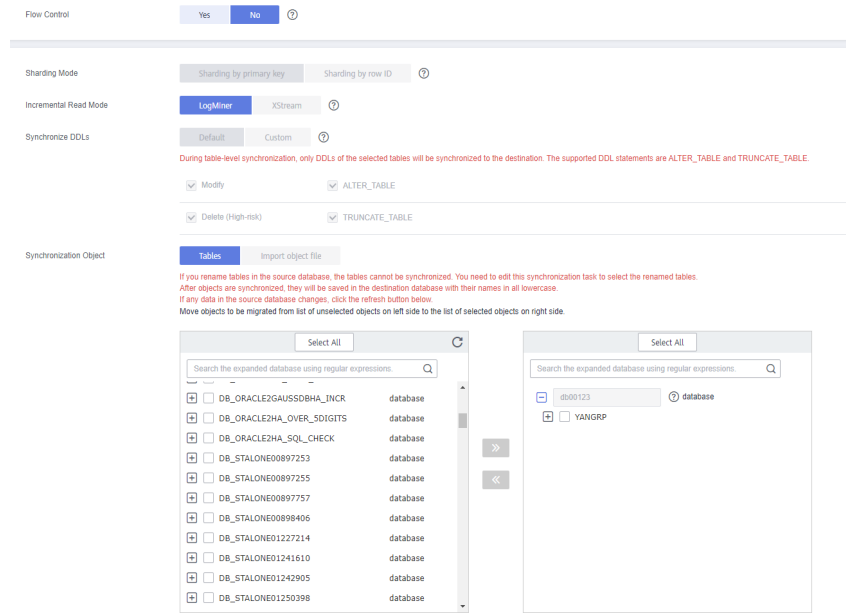
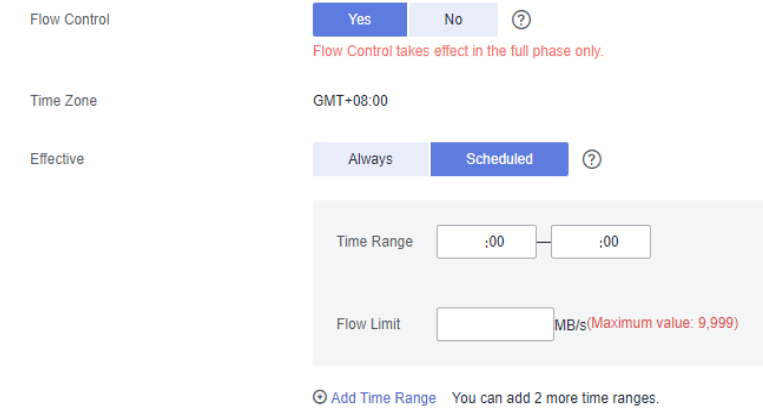



Table 3-170 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-160 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |


| Parameter | Description |
|-----------------------|---|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations. |
| Incremental Read Mode | <p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> • LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. • XStream: The XStream interface must be enabled. <p>LogMiner is recommended.</p> |
| Synchronize DDLs | <p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> • Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. • Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. |

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

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

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

Figure 3-161 Processing data



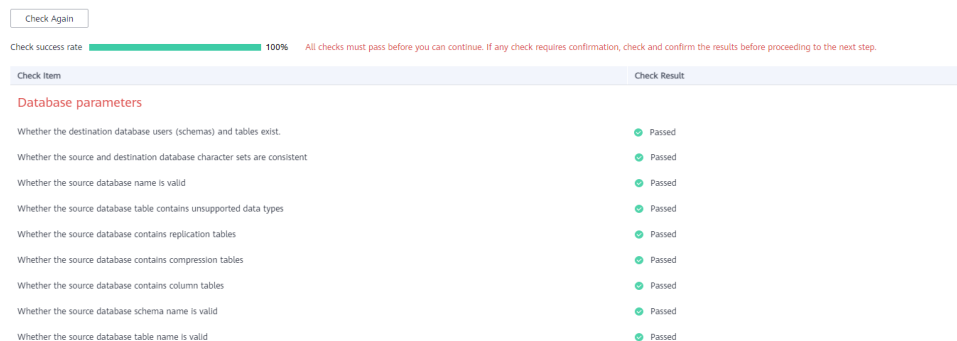
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-162 Pre-check



NOTE

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

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

Figure 3-163 Task startup settings

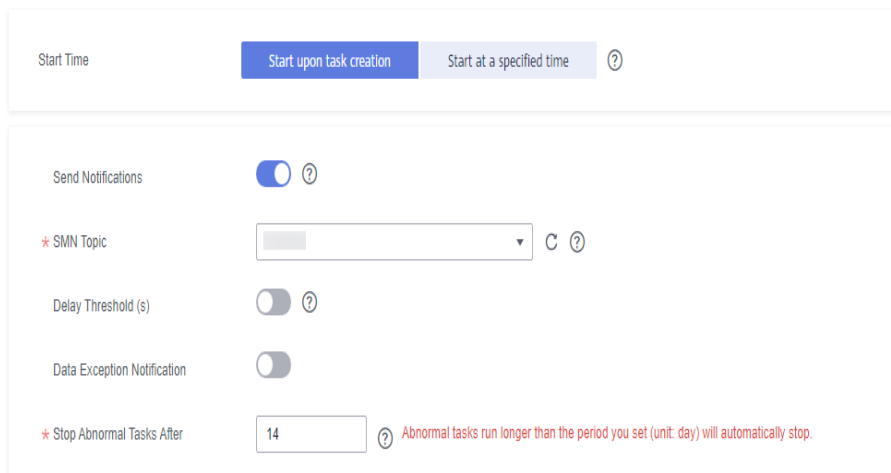



Table 3-171 Task startup settings

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

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

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

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

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

3.16 From Oracle to DDM

Supported Source and Destination Databases

Table 3-172 Supported databases

| Source DB | Destination DB |
|--|---|
| <ul style="list-style-type: none"> On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) | <ul style="list-style-type: none"> DDM instances |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-173 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 ON <userName.tbName> to drsUser). - Full+Incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|---|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"><li data-bbox="592 412 1412 510">• 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, urwid, interval (precision greater than 6 digits), and user-defined types are not supported. ● In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urwid, 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● 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. |

| Type | Restrictions |
|-----------------|--|
| | <ul style="list-style-type: none"> • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • Tables with virtual columns in the source database cannot be synchronized. |
| Source database | <ul style="list-style-type: none"> • The maximum row length of Oracle cannot exceed 8 KB, excluding BLOB and TEXT columns which will be automatically converted to the text and blob types of MySQL. The reason is that the MySQL InnoDB restricts the row length to 8 KB. • The primary key or unique key column cannot contain values of the string data type when you map the MySQL data types to the character data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur. • The values of binary_float and binary_double cannot be set to Nan, Inf, or -Inf because MySQL does not support these values. DRS converts the three values to 0 and saves them by default. • AUTO_PK_ROW_ID cannot be used as a column name in Oracle because it is a reserved column name in MySQL 5.7 and cannot be created. • Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle. • Tables whose names contain special characters except underscores (_) cannot be synchronized. • 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. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● 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 the synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● During 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. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. • During an incremental synchronization of table-level objects, renaming tables is not recommended. • The source database username is consistent with the destination database name. • When the source database user and table structure are synchronized to the destination database, the corresponding information is converted into lowercase letters. For example, tables Ab and AB are changed to table ab after being synchronized to the destination database. • Index organized tables cannot be synchronized. • Before a full synchronization task starts, if a transaction is not committed for a long time, data may be lost. • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. |

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 3-164 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 3-174 Task information

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

- Synchronization instance details

Figure 3-165 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, MariaDB, MongoDB, PostgreSQL, Microsoft SQL Server, GaussDB(for MySQL), TiDB

* Destination DB Engine: MySQL, **DDM**, GaussDB(DWS), GaussDB Distributed, GaussDB Primary/Standby, PostgreSQL, GaussDB(for MySQL)

* Network Type: Public network

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

Synchronization Instance Subnet: Select the subnet | View Subnets | View occupied IP address

* Synchronization Mode: **Full-incremental** | Full

Table 3-175 Synchronization instance information

| Parameter | Description |
|-----------|------------------------------|
| Data Flow | Select To the cloud . |

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

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE</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. 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. |

- Task type

Figure 3-166 Task type



Table 3-176 Task type information

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

- Enterprise Project and Tags

Figure 3-167 Enterprise projects and tags




Table 3-177 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-168 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 3-178 Source database settings

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

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

 **NOTE**


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

Figure 3-169 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 


 Test successful

Table 3-179 Destination database settings

| Parameter | Description |
|-------------------|--|
| DB Instance Name | The DDM instance you selected when you create a synchronization task. The instance name cannot be changed. |
| Database Username | The username for accessing the destination database. |
| Database Password | The database username and password are encrypted and stored in the system and will be cleared after the task is deleted. |

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

Figure 3-170 Synchronization Mode

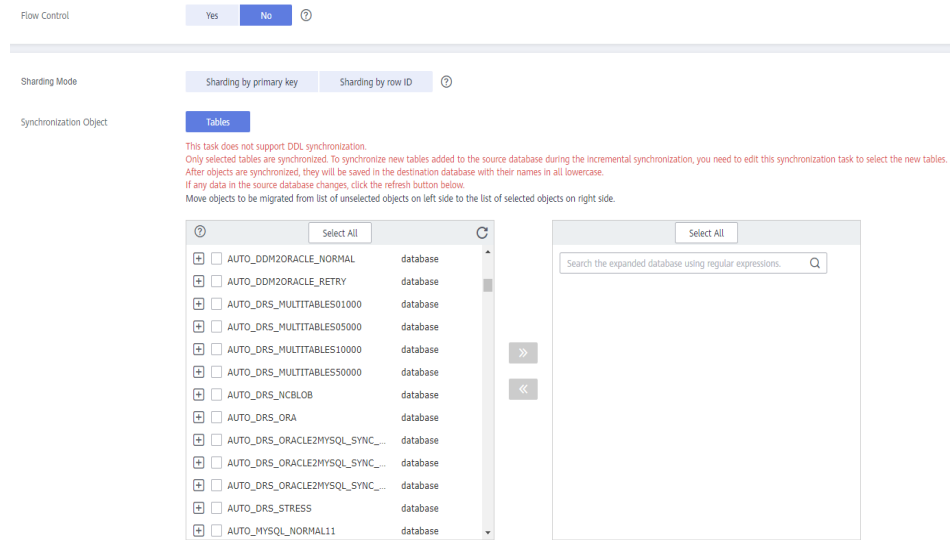
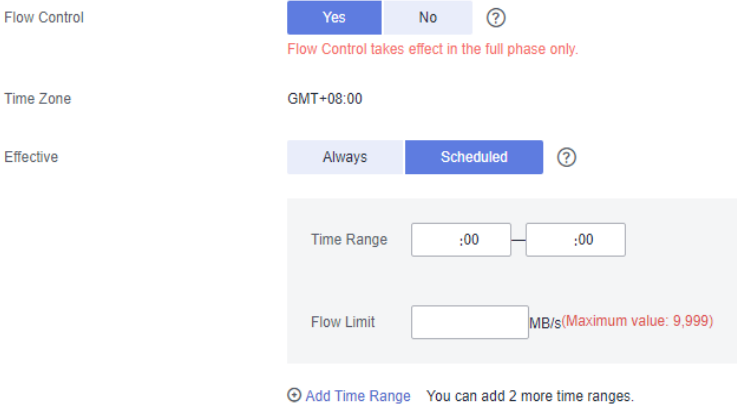



Table 3-180 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-171 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|------------------------|---|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations. |
| Synchronization Object | <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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

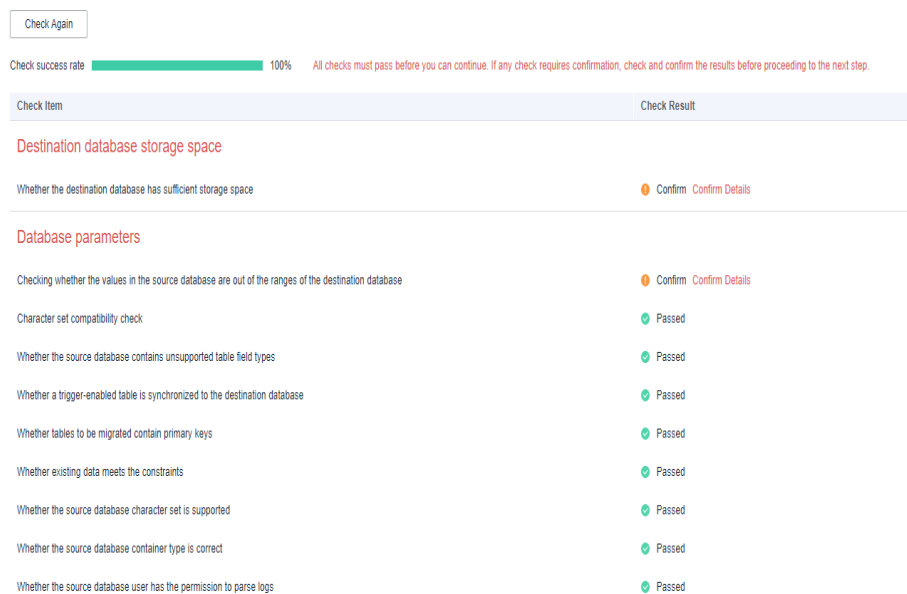
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-172 Pre-check



NOTE

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

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

Figure 3-173 Task startup settings

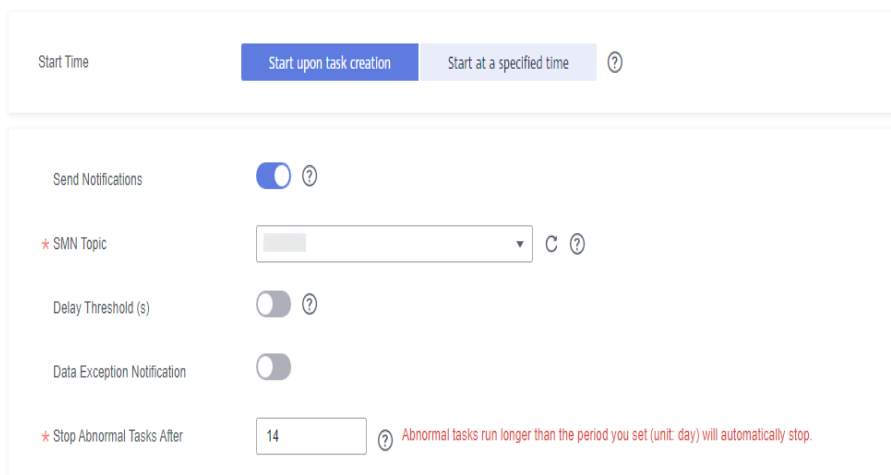



Table 3-181 Task startup settings

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

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

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

----End

3.17 From Oracle to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-182 Supported databases

| Source DB | Destination DB |
|--|--|
| <ul style="list-style-type: none"> • On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) • Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) | <ul style="list-style-type: none"> • GaussDB(DWS) cluster |

Suggestions

 **CAUTION**

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

-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-183 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 ON <userName.tbName> to drsUser). - Full+incremental synchronization and incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none">• 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 | <ul style="list-style-type: none"> ● Only tables, common indexes (B-Tree indexes), and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, unique indexes, foreign key indexes, and check constraints 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● 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. ● Only the to_date and sys_guid functions can be used as default values. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following situations may occur: <ul style="list-style-type: none"> - The default value is left empty. - The table fails to be created. As a result, the task fails. - If DDLs for incremental synchronization contain system functions that are not found in the destination database, the synchronization task fails. ● If 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 expression functions cannot be synchronized. |

| Type | Restrictions |
|----------------------|---|
| | <ul style="list-style-type: none"> • Temporary tables in the source database cannot be synchronized. • Virtual columns in the source database table cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. • When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. • Materialized views in the source Oracle database will be converted into common tables after being synchronized to the destination database. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. |
| 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. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● The 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 |

| 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. ● 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. ● During the synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● 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. |

| Type | Restrictions |
|------|---|
| | <p>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.</p> <ul style="list-style-type: none"> • When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. • During incremental synchronization, the following DDL operations are supported: DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database. <ul style="list-style-type: none"> - Table-level synchronization supports the following basic DDL operations: alter table add column, alter table drop column, alter table modify column, and truncate table. DDL operations such as create index are not supported. • The names of mapped databases and tables are case-insensitive, which means no matter if the object name is uppercase or lowercase, it stays lowercase after the object is synchronized to the destination database. • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. • You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental |

| Type | Restrictions |
|------|--|
| | data filtering conditions. Oracle logs may not record the old value of update . |

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 3-174 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-0131 [Text input]

Description: [Text area]

Table 3-184 Task information

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

- Synchronization instance details

Figure 3-175 Synchronization instance details

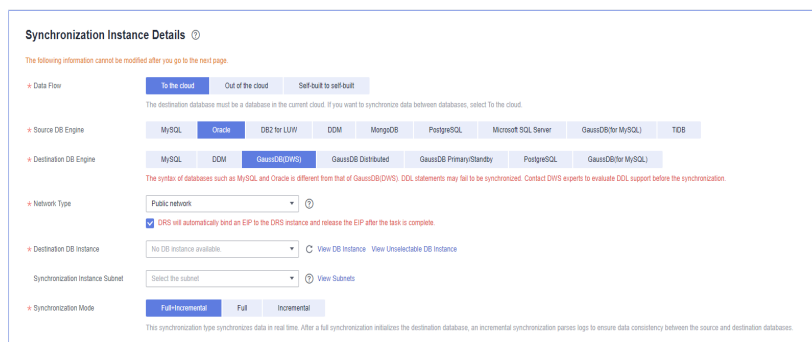


Table 3-185 Synchronization instance settings

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

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 3-176 Task type




Table 3-186 Task type information

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

- Enterprise Project and Tags

Figure 3-177 Enterprise projects and tags




Table 3-187 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-178 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 3-188 Source database settings

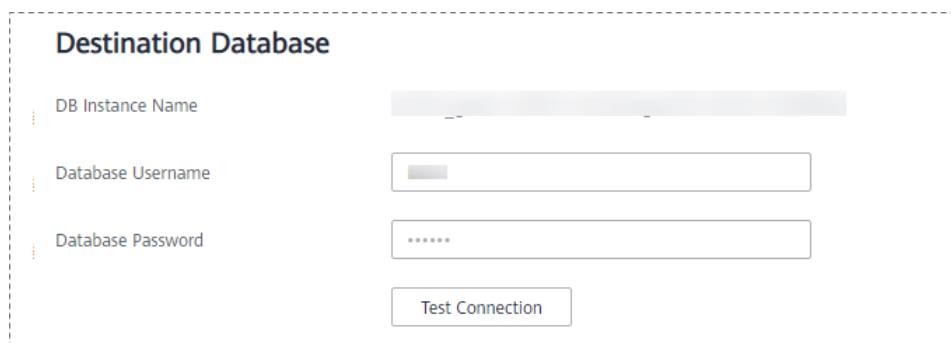
| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Service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

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

 **NOTE**

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

Figure 3-179 Destination database information



The screenshot shows a form titled "Destination Database" with three input fields: "DB Instance Name" (with a dropdown arrow), "Database Username" (with a small grey bar), and "Database Password" (with asterisks). A "Test Connection" button is located at the bottom right of the form.

Table 3-189 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 3-180 Synchronization mode

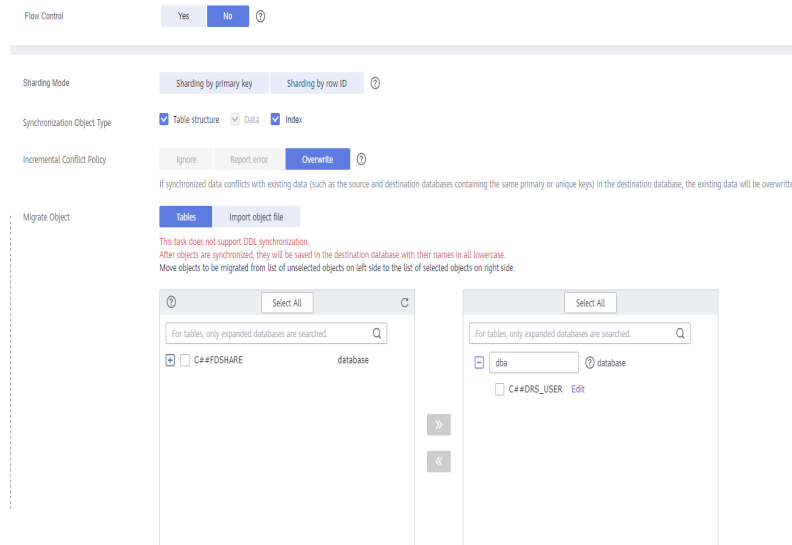
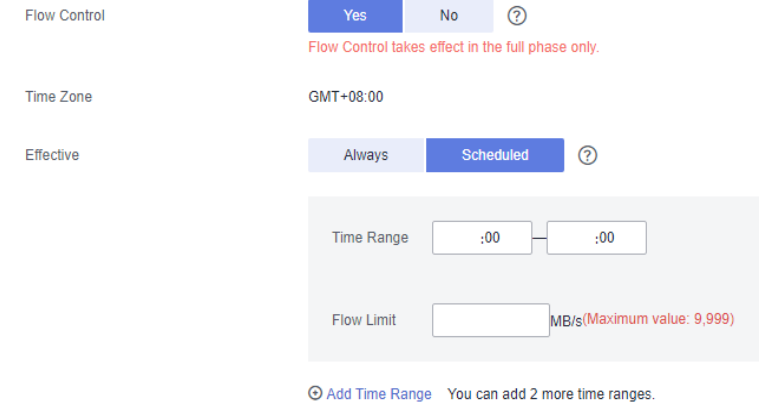



Table 3-190 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-181 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> ● Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. ● Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations. |
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> ● Data is selected by default. ● If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. ● If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. |

| Parameter | Description |
|------------------------|---|
| 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. <ul style="list-style-type: none"> In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

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

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

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

Figure 3-182 Processing data

| 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_MEASPUB_VSFYA | CHENHAO.IUFO_MEASPUB_VSFYA | -- | -- | -- | Add |
| CHENHAO.IUFO_MEASPUB_XSFYA | CHENHAO.IUFO_MEASPUB_XSFYA | -- | -- | -- | Add |
| CHENHAO.NA | CHENHAO.NA | -- | -- | -- | Add |
| CHENHAO.NB | CHENHAO.NB | -- | -- | -- | Add |
| CHENHAO.ONE | CHENHAO.ONE | -- | -- | -- | Add |
| CHENHAO.ONE1 | CHENHAO.ONE1 | -- | -- | -- | Add |
| CHENHAO.ORACLE1 | CHENHAO.ORACLE1 | -- | -- | -- | Add |

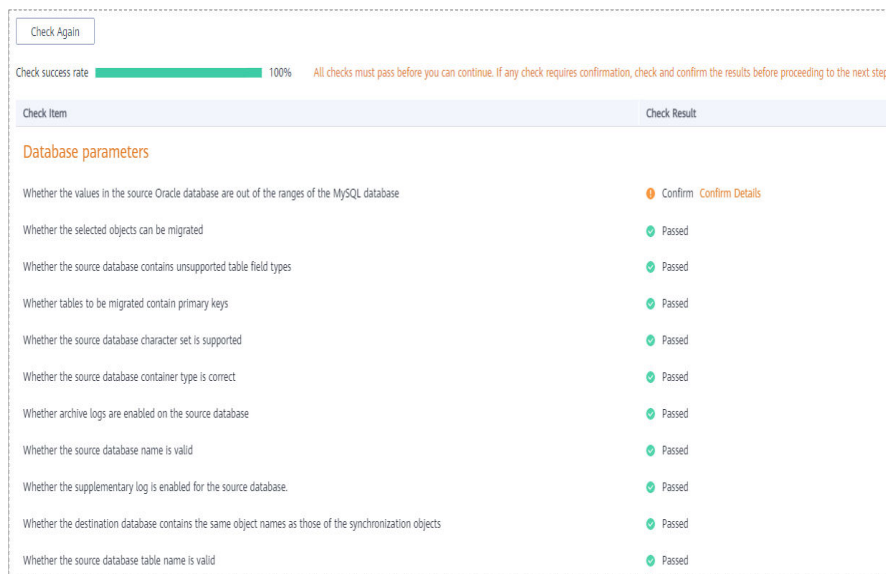
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-183 Pre-check



NOTE

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

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

Figure 3-184 Task startup settings

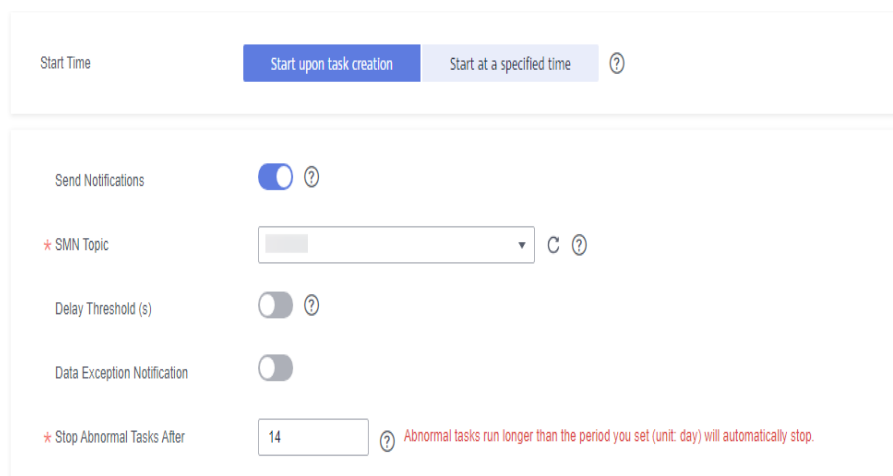



Table 3-191 Task startup settings

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

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

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

----End

3.18 From Oracle to PostgreSQL

Supported Source and Destination Databases

Table 3-192 Supported databases

| Source DB | Destination DB |
|--|--|
| <ul style="list-style-type: none"> • On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) • Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) | <ul style="list-style-type: none"> • 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.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-193 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 ON <userName.tbName> to drsUser). - Full+Incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"><li data-bbox="592 412 1369 508">• 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● 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 expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. |

| Type | Restrictions |
|----------------------|---|
| | <ul style="list-style-type: none"> • Tables with virtual columns 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 after time zone conversion. |
| 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 PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● 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. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> ● 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. ● 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. ● During the synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● 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. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental |

| Type | Restrictions |
|------|--|
| | <p>synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database.</p> <ul style="list-style-type: none"> ● 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, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● In a full+incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database. <ul style="list-style-type: none"> – Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. ● The 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 |

| Type | Restrictions |
|------|---|
| | <p>synchronization to reduce the number of WAL logs. After the synchronization is complete, enable the function. For details, see Setting an Automated Backup Policy.</p> <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</p> |

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 3-185 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 small note below the Region dropdown 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." The Task Name field has a character count of 0/256, and the Description field has a character count of 0/256.

Table 3-194 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 3-186 Synchronization instance details




Table 3-195 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE</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"> – 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. |

- Task type

Figure 3-187 Task type




Table 3-196 Task type information

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

- Enterprise Project and Tags

Figure 3-188 Enterprise projects and tags




Table 3-197 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-189 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 3-198 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | <p>The IP address or domain name of the source database.</p> <p>NOTE For a RAC cluster, use a Scan IP address to improve access performance.</p> |

| Parameter | Description |
|-----------------------|---|
| Port | The port of the source database. Range: 1 – 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| PDB Name | Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **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 3-190 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 3-199 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 3-191 Synchronization mode

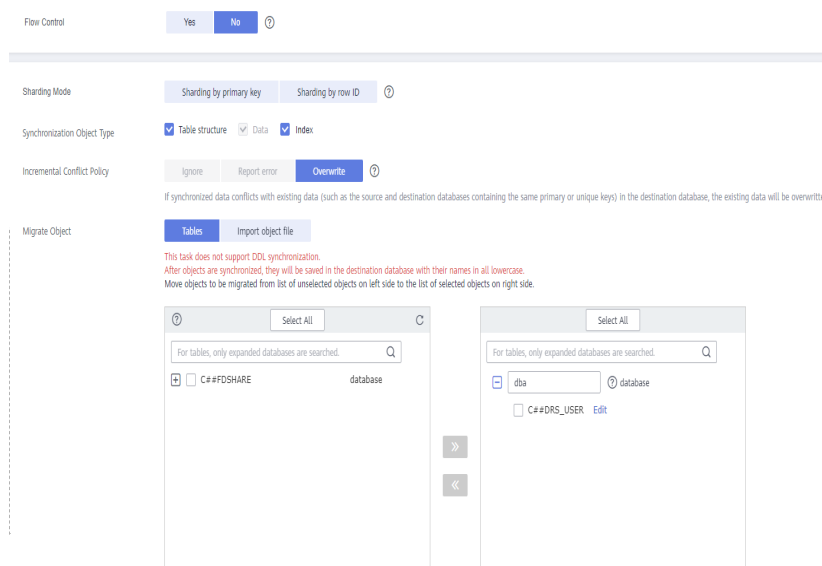
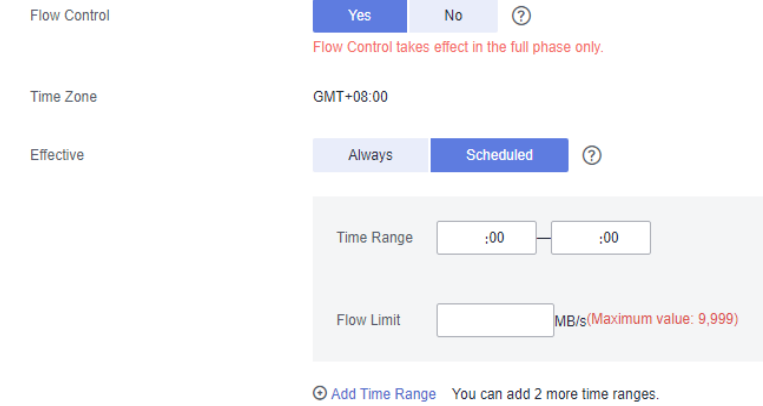



Table 3-200 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-192 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

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

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

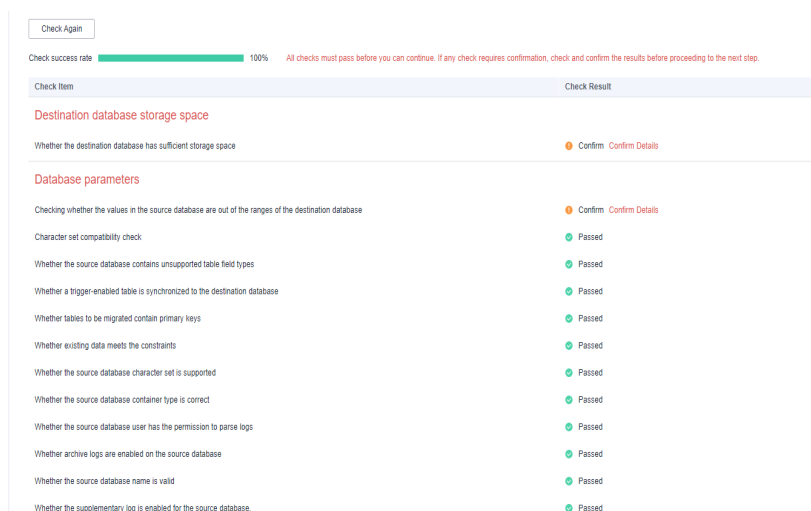
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 3-194 Task startup settings

Table 3-201 Task startup settings

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

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

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

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

----End

3.19 From DDM to MySQL

Supported Source and Destination Databases

Table 3-202 Supported databases

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

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-203 Precautions

| Type | Restrictions |
|------------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The destination database user must have the following permissions: 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. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● Resumable upload is supported, 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After a task is created, the destination database cannot be set to read-only. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. ● 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 full data synchronization, a lot of binlogs are generated in the destination database, occupying too much |

| Type | Restrictions |
|------|---|
| | <p>storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket.</p> <ul style="list-style-type: none"> ● 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. - Online DDL is not supported. - Tables can be created. For example: <pre>create table `ddl_test` (id int, c1 varchar(25), primary key(id)); create table `ddl_test_gho` like `ddl_test`;</pre> - Tables can be renamed. Both the source and destination tables must be selected. For example: <pre>rename table `ddl_test` to `ddl_test_new`;</pre> - Columns in a table can be added and modified, but cannot be deleted. For example: <pre>alter table `ddl_test` add column `c2` varchar(25); alter table `ddl_test` modify column `c1` varchar(50); alter table `ddl_test` alter c1 set default '***';</pre> - Table indexes can be modified. For example: <pre>alter table `ddl_test` drop primary key; alter table `ddl_test` add primary key(id); alter table `ddl_test` add index `ddl_test_uk`(id); alter table `ddl_test` drop index `ddl_test_uk`;</pre> - In table-level synchronization, you can add columns, modify columns, and add primary keys and normal indexes. - 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. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> When editing a task, do not import a lot of data to the newly-added table. You are advised to edit the task during off-peak hours. |

Procedure

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

Figure 3-195 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 text area with a height of 0/256 characters and a help icon.

Table 3-204 Task information

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

- Synchronization instance details

Figure 3-196 Synchronization instance details

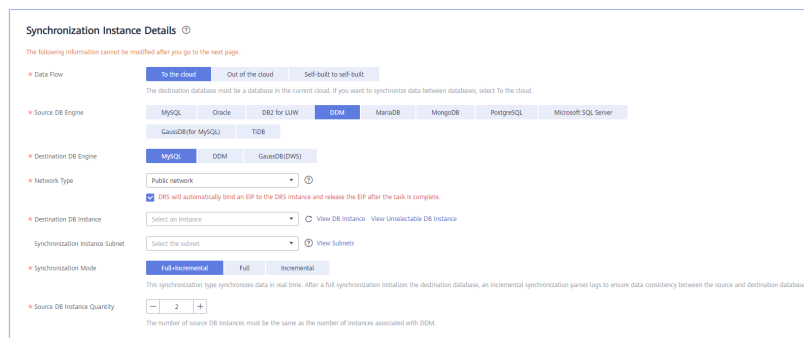


Table 3-205 Synchronization instance settings

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

| Parameter | Description |
|-----------------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |
| Source DB Instance Quantity | Specifies the number of DB instances bound to the source DDM database. The default value is 2 . The value ranges from 1 to 64 . Set this parameter based on the site requirements. |

- Task type

Figure 3-197 Task type



Table 3-206 Task type information

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

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

- Enterprise Project and Tags

Figure 3-198 Enterprise projects and tags




Table 3-207 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-199 Source database information

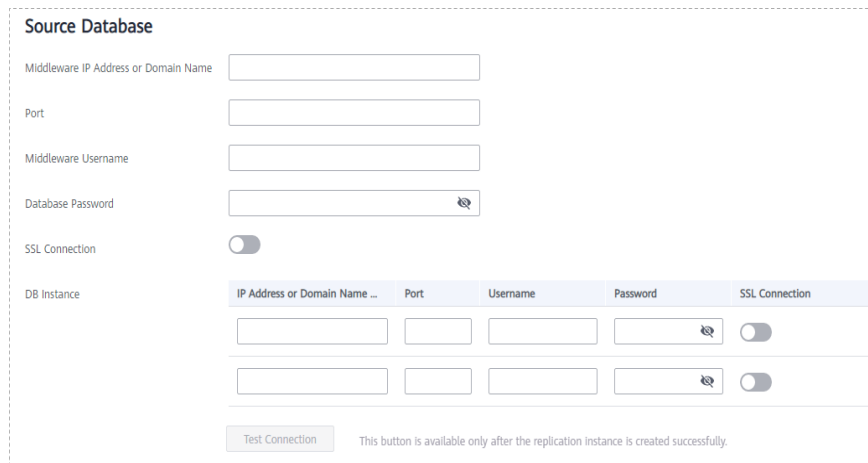


Table 3-208 Source database settings

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

NOTE


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

Figure 3-200 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-209 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 disabled, your data may be at risk. |

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

Figure 3-201 Synchronization mode

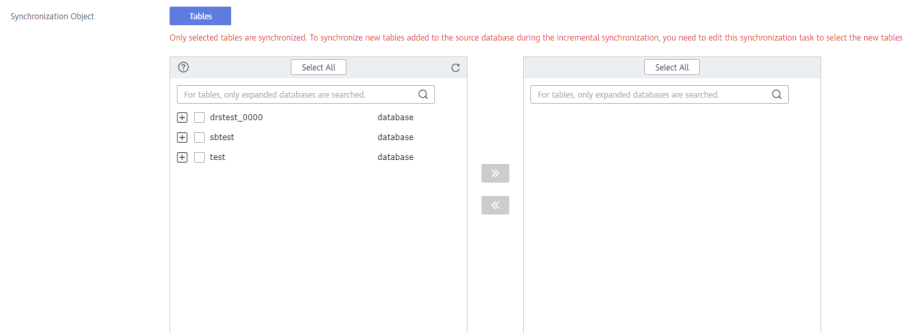



Table 3-210 Synchronization mode and object

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

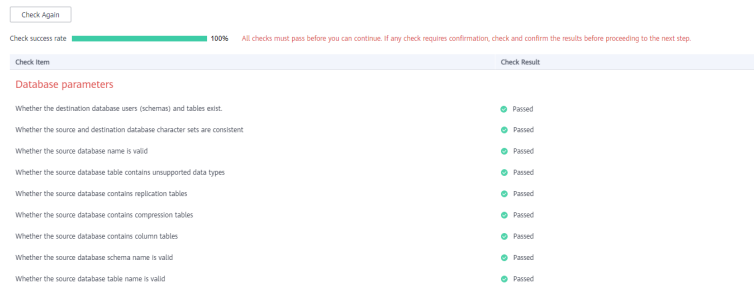
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-202 Pre-check



NOTE

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

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

Figure 3-203 Task startup settings

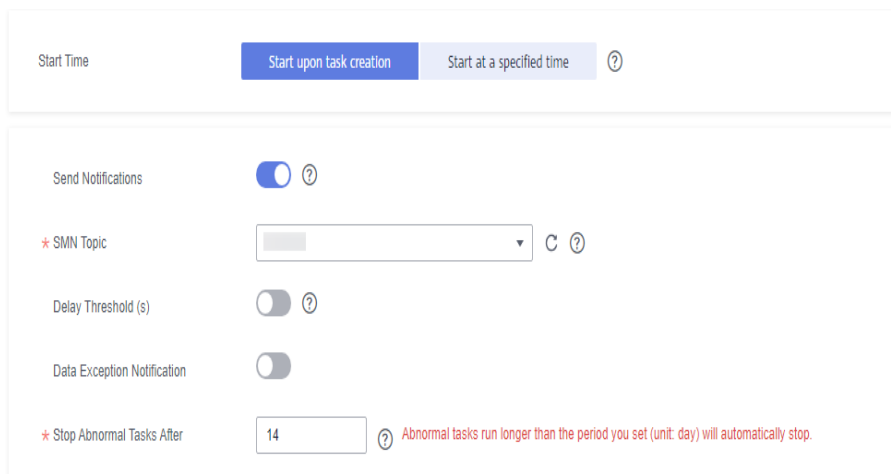



Table 3-211 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |

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

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

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

----End

3.20 From DDM to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-212 Supported databases

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

Suggestions

⚠ CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-213 Precautions

| Type | Restrictions |
|------------------------|--|
| Database permissions | <ul style="list-style-type: none"> • The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE. |
| Synchronization object | <ul style="list-style-type: none"> • Full synchronization supports the synchronization of data, table structures, and common indexes (B-Tree indexes). • The destination database does not support tables with unique keys. During synchronization, the table with unique keys in the source database is ignored. • A table without a primary key cannot be synchronized. If the selected table does not have a primary key, the synchronization fails. • Unique constraints and indexes cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. |

| 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. ● In the source database, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB(DWS), the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: hash value + original constraint name (which may be truncated) + _key, or table name_original index name. ● In the full synchronization phase, only B-Tree indexes are synchronized. Other indexes are not synchronized by default. If there are too many GaussDB(DWS) indexes, the storage space and data import performance will be affected. You are advised to create indexes based on service requirements. ● 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. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • During an incremental synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During an incremental synchronization, if you need to modify the structure of the source table to be synchronized, you must modify the corresponding destination table structure. • During 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. • During an incremental synchronization, tables whose primary key type is binary, text, blob, or clob cannot be deleted or updated. • When editing a task, do not import a lot of data to the newly-added table. You are advised to edit the task during off-peak hours. |

Procedure

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

Figure 3-204 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. 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 includes:

- Region:** A dropdown menu with a downward arrow.
- Project:** A dropdown menu with a downward arrow.
- * Task Name:** A text input field containing "DRS-8/131" and a help icon (question mark in a circle).
- Description:** A larger text input area with a help icon (question mark in a circle).

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

Table 3-214 Task information

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

- Synchronization instance details

Figure 3-205 Synchronization instance details

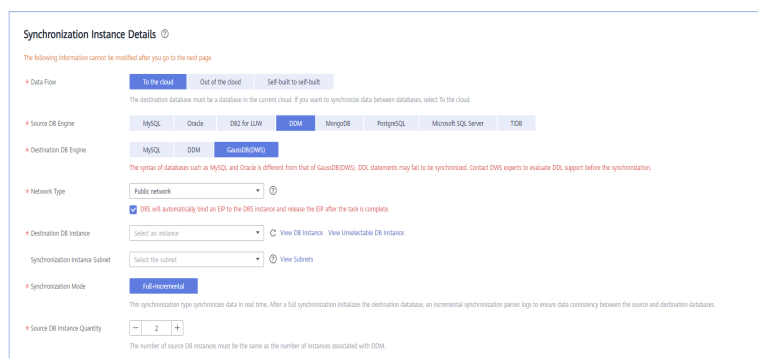


Table 3-215 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|---|
| Network Type | <p>Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | An available GaussDB(DWS) instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |
| Source DB Instance Quantity | Specifies the number of DB instances bound to the source DDM database. The default value is 2 . The value ranges from 1 to 64 . Set this parameter based on the site requirements. |

- Task type

Figure 3-206 Task type



Table 3-216 Task type information

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

- Enterprise Project and Tags

Figure 3-207 Enterprise projects and tags

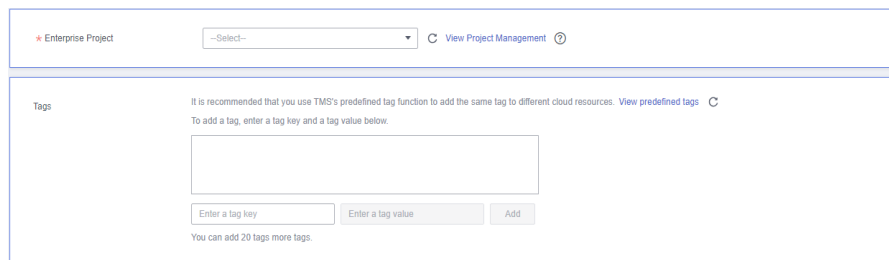


Table 3-217 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-208 Source database information

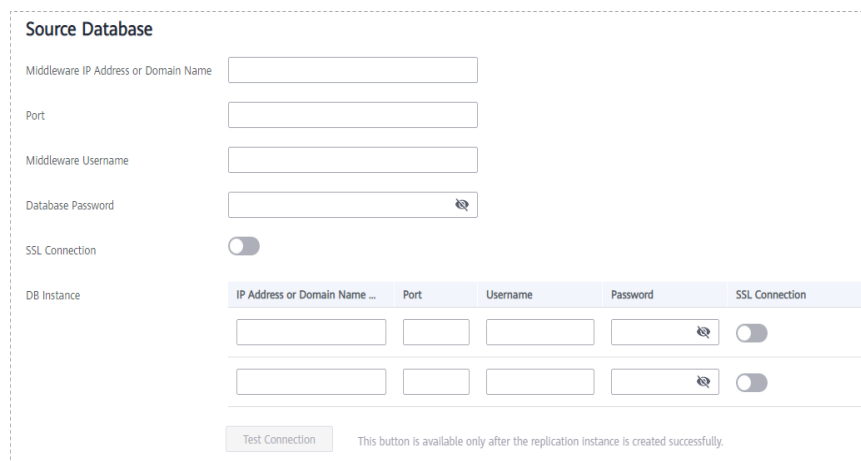


Table 3-218 Source database settings

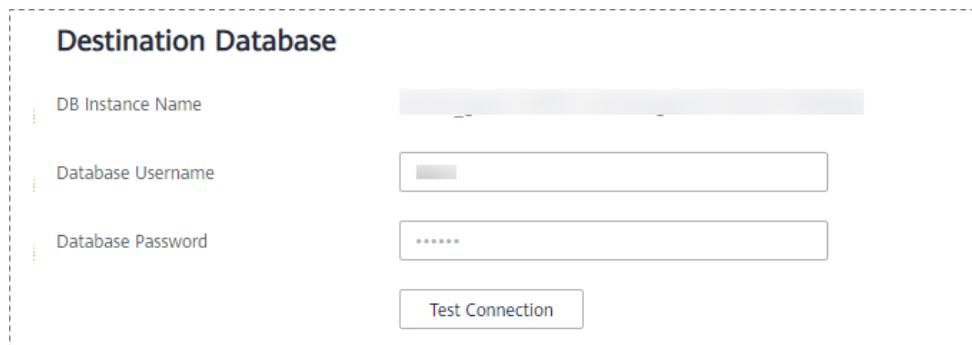
| Parameter | Description |
|--------------------------------------|---|
| Middleware IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |

| Parameter | Description |
|---------------------|--|
| Middleware Username | The username of the source DDM instance. |
| Middleware Password | The password for the source DDM instance username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |
| DB Instance | The sharded database details. |

 **NOTE**

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

Figure 3-209 Destination database information



Destination Database

DB Instance Name

Database Username

Database Password

Table 3-219 Destination database settings

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

NOTE

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

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

Figure 3-210 Synchronization mode

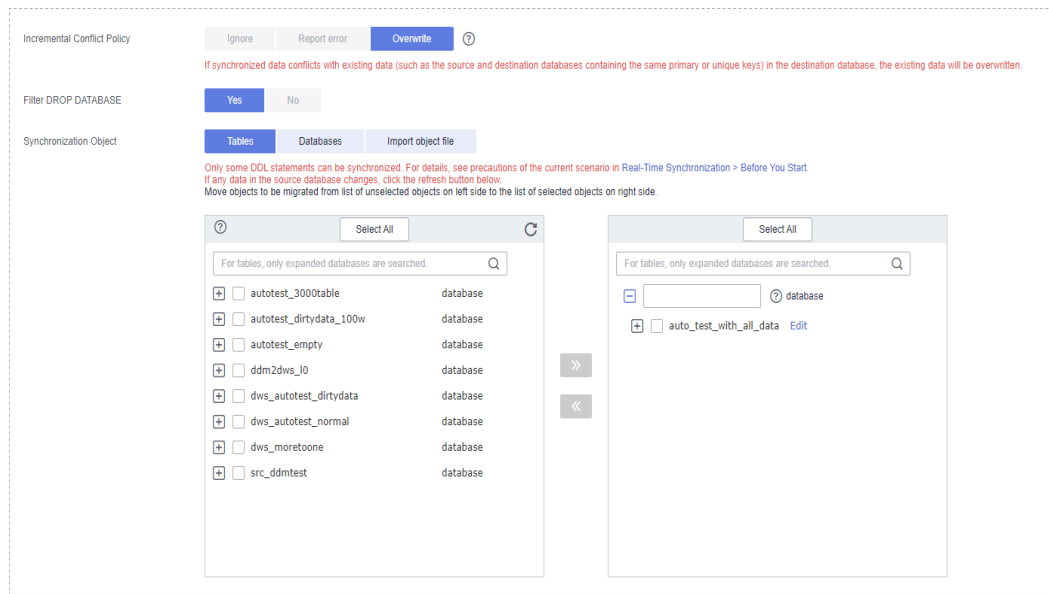




Table 3-220 Synchronization mode and object

| 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. <ul style="list-style-type: none"> In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Data Processing** page, select the table object to be processed, enter the column name, type, and operation type to be added, confirm the information, and click **Next**. You can configure related rules by referring to [Processing Data](#).

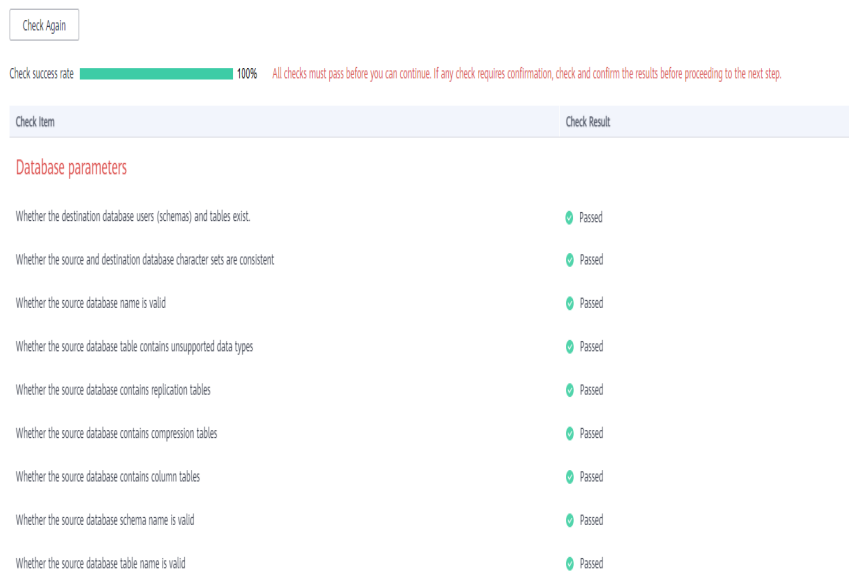
Figure 3-211 Processing data



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

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 3-212 Pre-check



NOTE

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

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

Figure 3-213 Task startup settings





Table 3-221 Task startup settings

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

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

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

----End

3.21 From DDM to DDM

Supported Source and Destination Databases

Table 3-222 Supported databases

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

Suggestions

 **CAUTION**

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

-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-223 Precautions

| Type | Restrictions |
|------------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The 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. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. ● 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. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> When editing a task, do not import a lot of data to the newly-added table. You are advised to edit the task during off-peak hours. If the destination DDM version is later than 3.0.4.1, DRS automatically updates the start value of the DDM sequence when the task is complete. 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 3-214 Synchronization task information

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

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 [Info icon]

Description: [Text area] [Info icon]

0/256

Table 3-224 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

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

- Synchronization instance details

Figure 3-215 Synchronization instance details

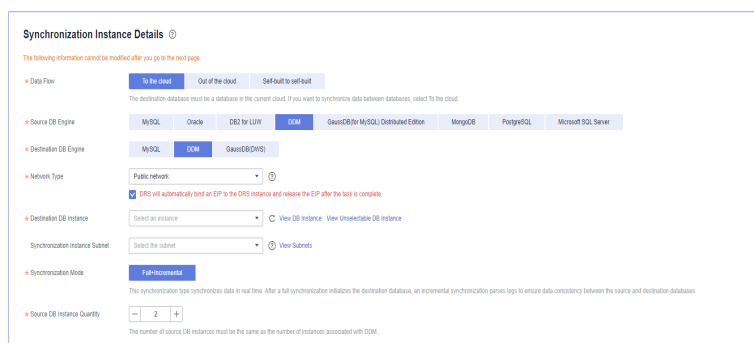


Table 3-225 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |
| Source DB Instance Quantity | <p>Specifies the number of DB instances bound to the source DDM database. The default value is 2. The value ranges from 1 to 64. Set this parameter based on the site requirements.</p> |

- Task type

Figure 3-216 Task type



Table 3-226 Task type information

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

- Enterprise Project and Tags

Figure 3-217 Enterprise projects and tags




Table 3-227 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-218 Source database information

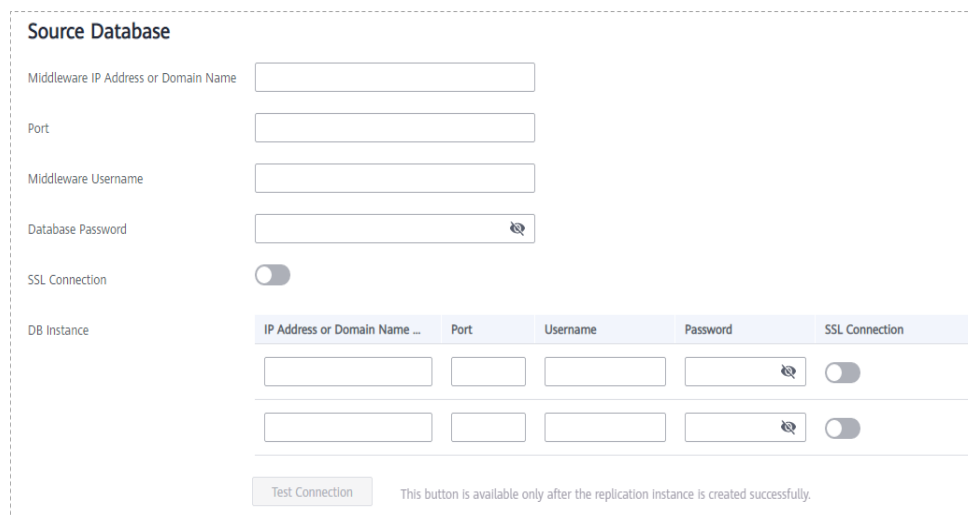


Table 3-228 Source database settings

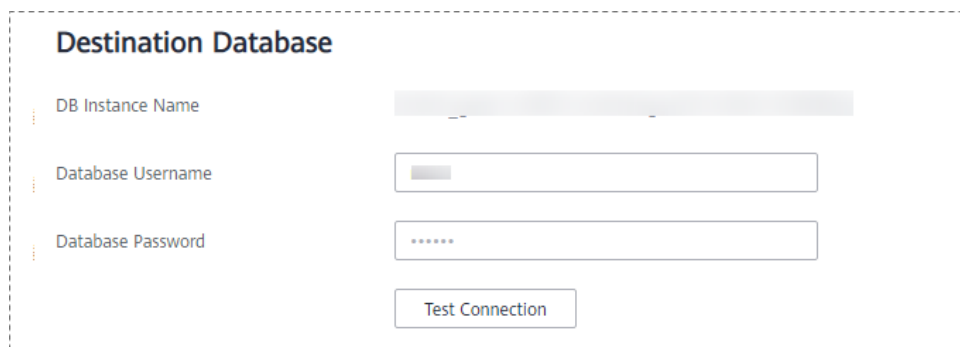
| Parameter | Description |
|--------------------------------------|---|
| Middleware IP Address or Domain Name | The IP address or domain name of the source DDM middleware. |

| Parameter | Description |
|---------------------|--|
| Port | The port of the source DDM middleware. Value range: 1 to 65535 |
| Middleware Username | The username of the source DDM instance. |
| Middleware Password | The password for the source DDM instance username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |
| DB Instance | The sharded database details. |

 **NOTE**

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

Figure 3-219 Destination database information



The screenshot shows a configuration form 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 asterisks). Below these fields is a "Test Connection" button.

Table 3-229 Destination database settings

| Parameter | Description |
|-------------------|--|
| DB Instance Name | The DDM instance you selected when you create the synchronization task. The instance name cannot be changed. |
| Database Username | The username for accessing the destination database. |
| Database Password | The database username and password are encrypted and stored in the system and will be cleared after the task is deleted. |

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

Figure 3-220 Synchronization mode

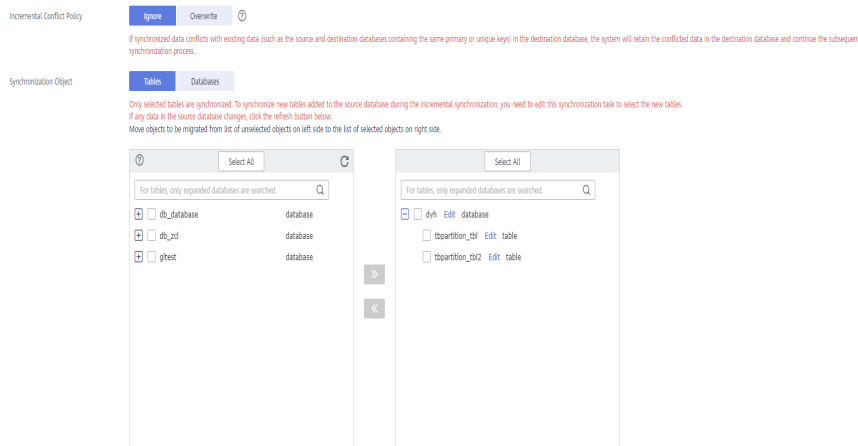



Table 3-230 Synchronization mode and object

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

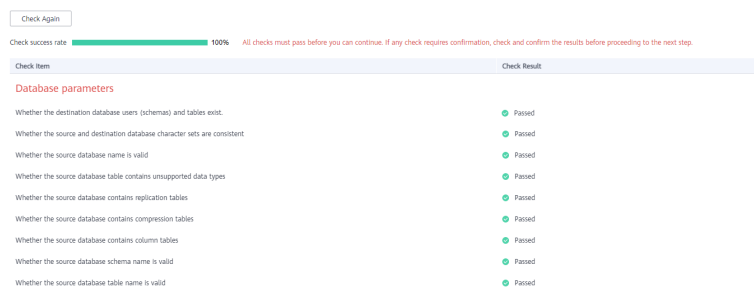
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-221 Pre-check



NOTE

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

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

Figure 3-222 Task startup settings

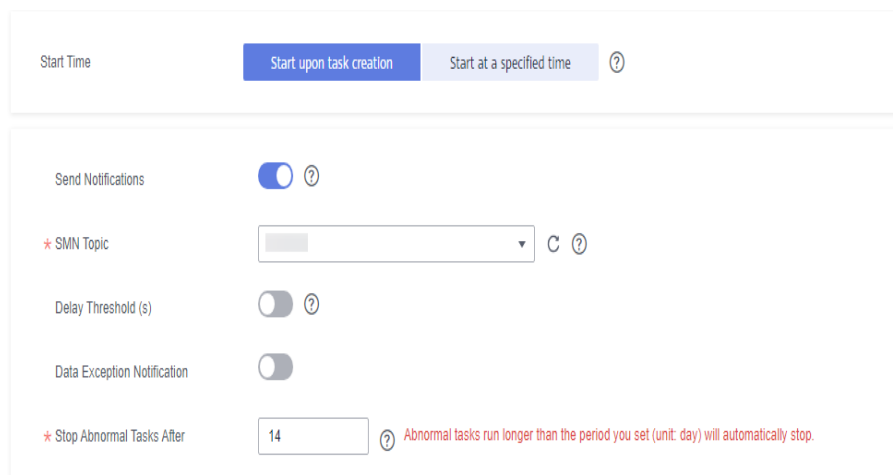



Table 3-231 Task startup settings

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

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

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

----End

3.22 From DB2 for LUW to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-232 Supported databases

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

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-233 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 permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. - Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. - Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary 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 maximum precision supported by the timestamp type is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. ● If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency. |
| Source database | <ul style="list-style-type: none"> ● Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or special characters .>`<`\, ?!". Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. ● Only the GBK and UTF8 character sets can be synchronized. |

| Type | Restrictions |
|----------------------|--|
| Destination database | <ul style="list-style-type: none">● The destination database must be a GaussDB primary/standby instance 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. ● After a task is started, the detach operation on a partition table may cause data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot contain uncommitted DDL transactions. ● Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. ● During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. ● Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • You can add additional objects in the incremental synchronization phase. |

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 3-223 Synchronization task information

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

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 [Text input]

Description: [Text area]

0/256

Table 3-234 Task information

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

- Synchronization instance details

Figure 3-224 Synchronization instance details

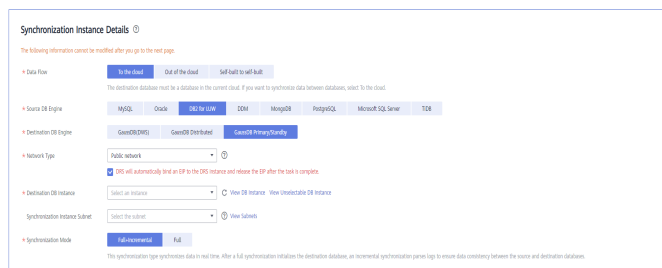


Table 3-235 Synchronization instance settings

| Parameter | Description |
|---------------------------------|--|
| Data Flow | Select To the cloud . |
| Source DB Engine | Select DB2 for LUW. |
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | <p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . |

- AZ

Figure 3-225 AZ



Table 3-236 Task AZ

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

- Enterprise Project and Tags

Figure 3-226 Enterprise projects and tags

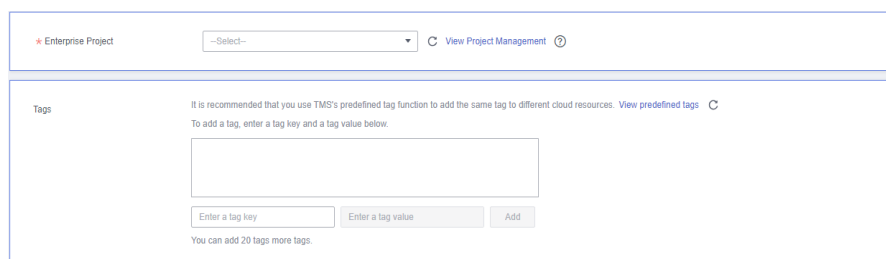


Table 3-237 Enterprise Project and Tags

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

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

 **NOTE**

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

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

Figure 3-227 Source database information


Source Database

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

IP Address or Domain Name

Port

Database Username

Database Password 

Database Name

SSL Connection

Synchronize Driver Synchronized.jcc-11.5.6.0(2).jar

Table 3-238 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. |

| Parameter | Description |
|--------------------|--|
| Database Password | The password for the source database username. |
| Database Name | The name of the database to which the synchronization objects belong in the source DB2 for LUW. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If the SSL certificate is not used, your data may be at risk. Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures. |
| Synchronize Driver | You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection. |

 **NOTE**

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

Figure 3-228 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-239 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 3-229 Synchronization mode

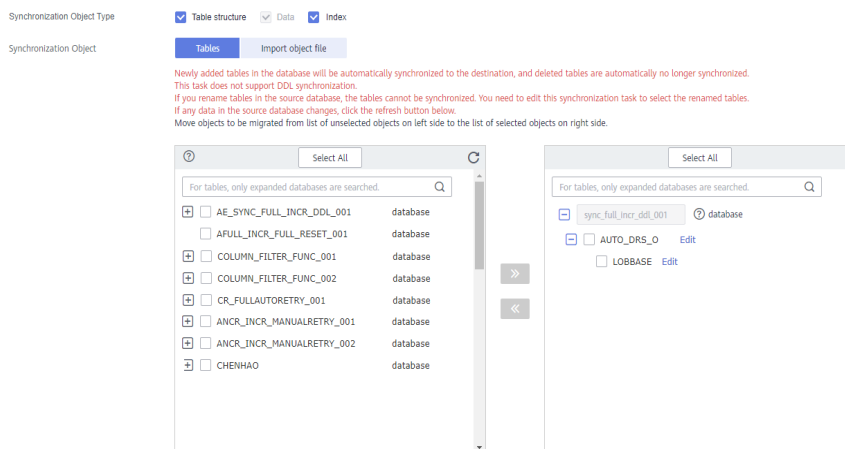



Table 3-240 Synchronization Object

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

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

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

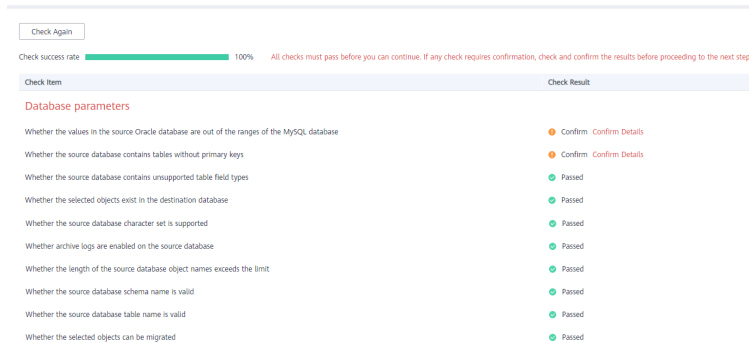
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-230 Pre-check



NOTE

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

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

Figure 3-231 Task startup settings

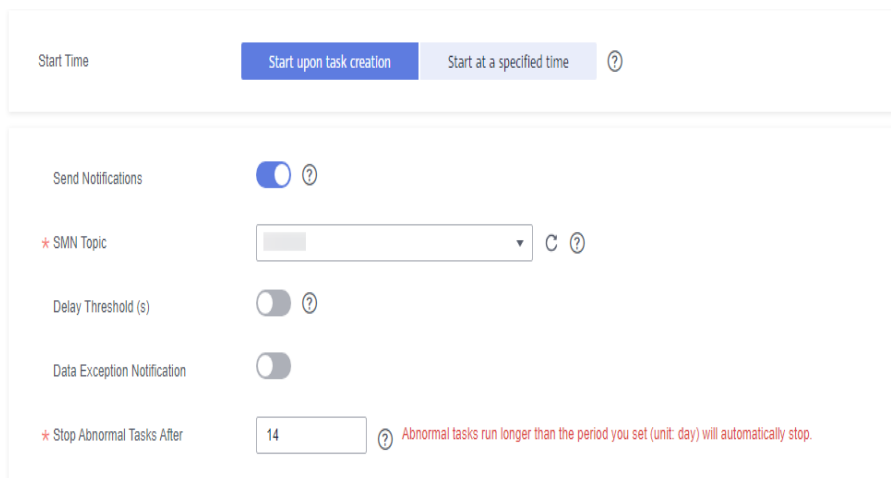



Table 3-241 Task startup settings

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

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

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

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

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

3.23 From DB2 for LUW to GaussDB Distributed

Supported Source and Destination Databases

Table 3-242 Supported databases

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

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-243 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 permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. - Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. - Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary 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 maximum precision supported by the timestamp type is 6. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • DDL statements executed in the source database cannot be synchronized. • For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. • If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency. |
| Source database | <ul style="list-style-type: none"> • Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or special characters: .>`<'\, ?!" Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. • Only the GBK and UTF8 character sets can be synchronized. |

| Type | Restrictions |
|----------------------|--|
| Destination database | <ul style="list-style-type: none"> ● The destination database must be a GaussDB distributed instance 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. ● After a task is started, the detach operation on a partition table may cause data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot contain uncommitted DDL transactions. ● Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. ● During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. ● Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> • During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. • You can add additional objects in the incremental synchronization phase. • During an incremental synchronization, if you update the primary key column or the first column of a table that does not have a primary key in 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 3-232 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 3-244 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 3-233 Synchronization instance details

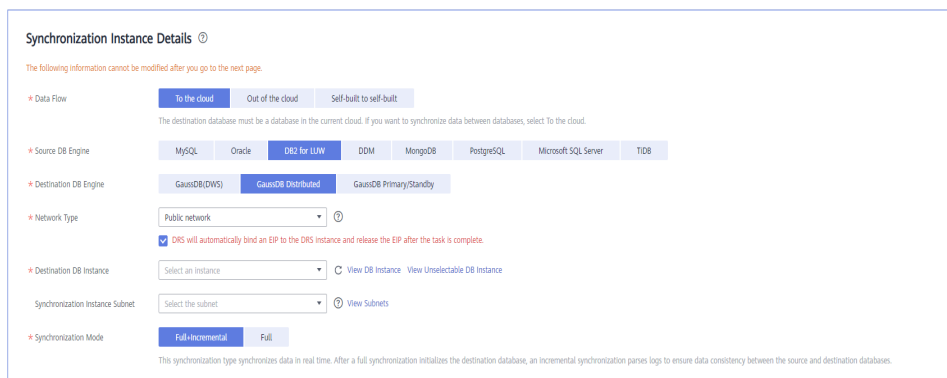


Table 3-245 Synchronization instance settings

| Parameter | Description |
|-------------------------|--|
| Data Flow | Select To the cloud . |
| Source DB Engine | Select DB2 for LUW . |
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | <p>Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | The GaussDB distributed instance you created. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Synchronization Mode | The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . |

- AZ

Figure 3-234 AZ




Table 3-246 Task AZ

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

- Enterprise Project and Tags

Figure 3-235 Enterprise projects and tags

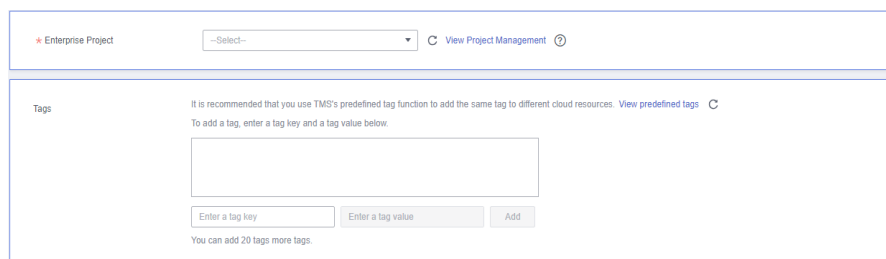


Table 3-247 Enterprise Project and Tags

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

 **NOTE**

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

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

Figure 3-236 Source database information

Source Database

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

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronized:jcc-11.5.6.0(2).jar

Table 3-248 Source database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | IP address or domain name of the source DB2 for LUW database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source DB2 for LUW database. |
| Database Password | The password for the source database username. |
| Database Name | The name of the database to which the synchronization objects belong in the source DB2 for LUW. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> • The maximum size of a single certificate file that can be uploaded is 500 KB. • If the SSL certificate is not used, your data may be at risk. • Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth • DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures. |
| Synchronize Driver | You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection. |

 **NOTE**

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

Figure 3-237 Destination database information

The screenshot shows a form titled "Destination Database" with three input fields and a button. The fields are labeled "DB Instance Name", "Database Username", and "Database Password". The "DB Instance Name" field is a wide text box. The "Database Username" and "Database Password" fields are narrower text boxes. Below the fields is a button labeled "Test Connection".

Table 3-249 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 3-238 Synchronization mode

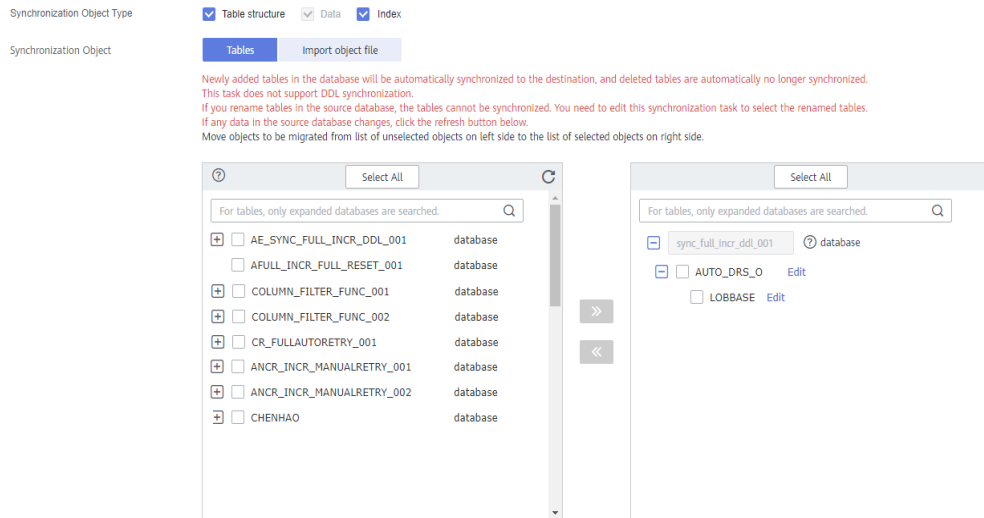



Table 3-250 Synchronization Object

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |

| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

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

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

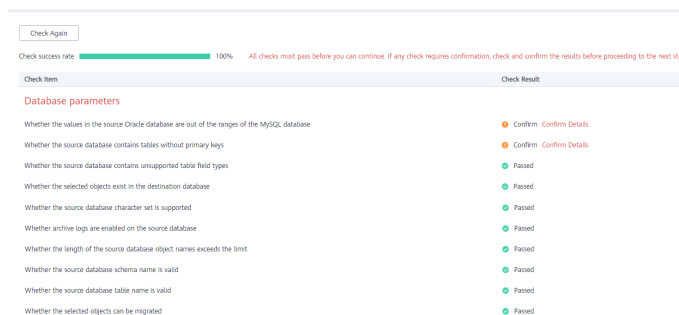
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-239 Pre-check



 **NOTE**

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

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


Figure 3-240 Task startup settings

Table 3-251 Task startup settings

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

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

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

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

----End

3.24 From DB2 for LUW to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-252 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) |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-253 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 maximum precision supported by the timestamp type is 6. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • DDL statements executed in the source database cannot be synchronized. • For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. • 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. • If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency. • The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. |
| 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.• After a task is started, the detach operation on a partition table may cause data inconsistency.• During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• During the synchronization, the source database cannot contain uncommitted DDL transactions.• Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database.• When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table.• During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. You can add additional objects in the incremental synchronization phase. |

Procedure

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

Figure 3-241 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 close button (X). Below the banner, the form contains the following fields:

- Region:** A dropdown menu with a downward arrow. Below it, a note 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."
- Project:** A dropdown menu with a downward arrow.
- Task Name:** A text input field containing "DRS-6131" and a help icon (i).
- Description:** A text area with a height of 10256 characters and a help icon (i).

Table 3-254 Task information

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

- Synchronization instance information

Figure 3-242 Synchronization instance information

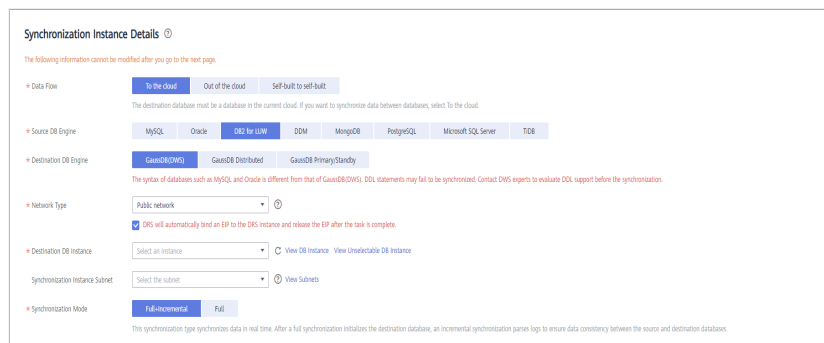


Table 3-255 Synchronization instance settings

| Parameter | Description |
|-------------------------|---|
| Data Flow | Select To the cloud . The destination database is Databases on the current cloud . |
| Source DB Engine | Select DB2 for LUW . |
| Destination DB Engine | Select GaussDB(DWS) . |
| Network Type | Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | The GaussDB(DWS) instance you created. |

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

- AZ

Figure 3-243 AZ



Table 3-256 Task AZ

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

- Enterprise Project and Tags

Figure 3-244 Enterprise projects and tags




Table 3-257 Enterprise Project and Tags

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

NOTE

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

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

Figure 3-245 Source database information

Source Database

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

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronized:jcc-11.5.6.0(2).jar

Table 3-258 Source database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | IP address or domain name of the source DB2 for LUW database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source DB2 for LUW database. |
| Database Password | The password for the source database username. |
| Database Name | The name of the database to which the synchronization objects belong in the source DB2 for LUW. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If the SSL certificate is not used, your data may be at risk. Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures. |
| Synchronize Driver | You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection. |

 **NOTE**


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

Figure 3-246 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 


 Test successful

Table 3-259 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 3-247 Synchronization Mode

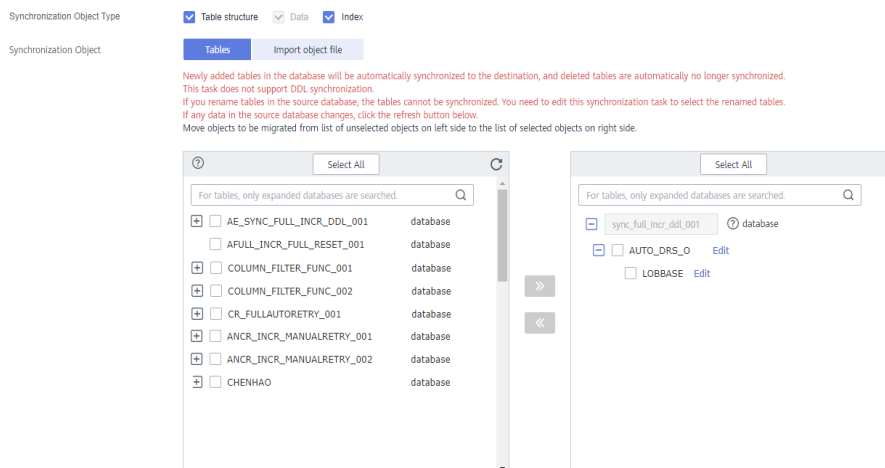



Table 3-260 Synchronization object

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

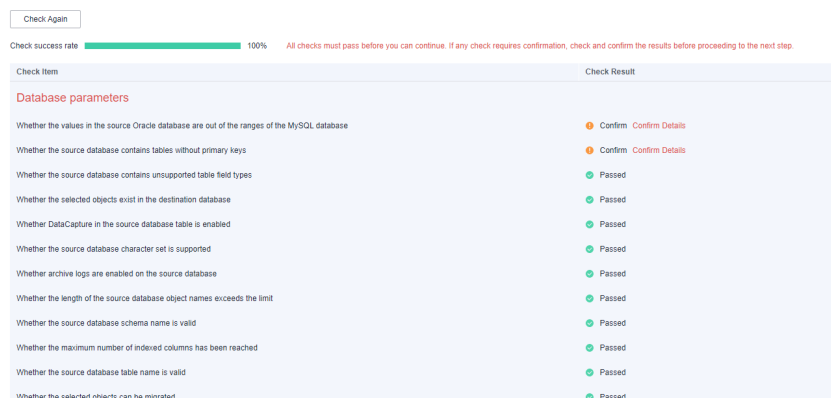
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-248 Pre-check



NOTE

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

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

Figure 3-249 Task startup settings

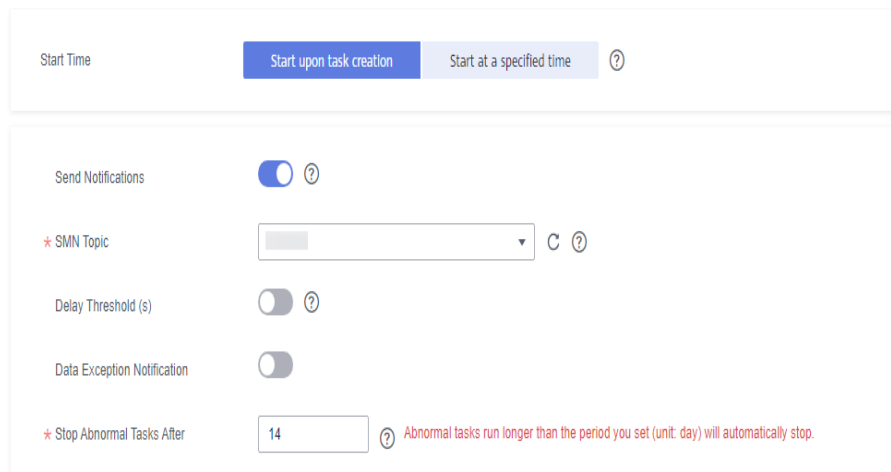



Table 3-261 Task startup settings

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

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

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

----End

3.25 From TiDB to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-262 Supported databases

| Source DB | Destination DB |
|--|------------------------------------|
| TiDB 4.0.0 and later (excluding the development version) | GaussDB(for MySQL) Primary/Standby |

Supported Synchronization Objects

[Table 3-263](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-263 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, 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 3-264](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

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

Table 3-264 Database account permission

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

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 3-265 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 <p>CAUTION</p> <ul style="list-style-type: none"> ● TiDB Binlog is incompatible with the following features in TiDB 5.0: <ul style="list-style-type: none"> TiDB clustered index: After TiDB Binlog is enabled, you cannot create a clustered index that is defined as a single non-integer primary key. The insert, delete, and update operations on existing tables with clustered indexes are not synchronized to the downstream applications through TiDB Binlog. To synchronize a table with clustered indexes, upgrade the TiDB version to 5.1. TiDB system variable <code>tidb_enable_async_commit</code>: After TiDB Binlog is enabled, the performance cannot be improved if this option is enabled. TiDB system variable <code>tidb_enable_1pc</code>: After TiDB Binlog is enabled, the performance cannot be improved if this option is enabled. ● TiDB Binlog is incompatible with the following features in TiDB 4.0.7: <ul style="list-style-type: none"> TiDB system variable <code>tidb_enable_amend_pessimistic_txn</code>: After TiDB Binlog is enabled, the data replicated from the TiDB Binlog may be inconsistent. <ul style="list-style-type: none"> ● 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 <code>lower_case_table_names</code> of the destination database is set to <code>1</code>, databases or tables containing uppercase letters cannot be synchronized. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database must be a primary/standby GaussDB(for MySQL) instance. - In full+incremental synchronization, if the destination database already has a table structure, ensure that the table structure is the same as that on the source database or contains columns in the source database. - The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Constraints |
|-----------------------------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Triggers enabled in the destination database cannot be associated with synchronization tables. <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - You can add additional objects during an incremental synchronization. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. |
| 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. |

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 3-250 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is a warning message in an orange box: "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 3-266 Task information

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

- Synchronization instance details

Figure 3-251 Synchronization instance details

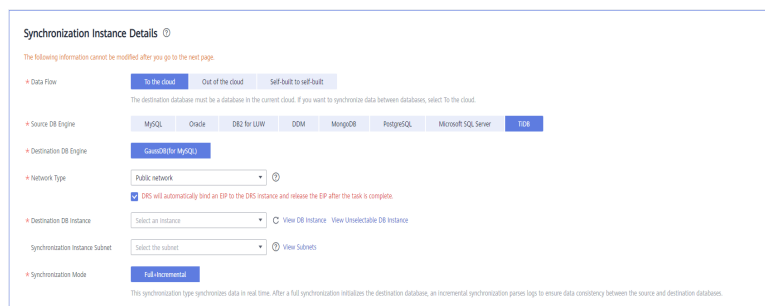


Table 3-267 Synchronization instance settings

| 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. <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| 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 Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. |

- AZ

Figure 3-252 AZ




Table 3-268 Task AZ

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

- Enterprise Project and Tags

Figure 3-253 Enterprise projects and tags

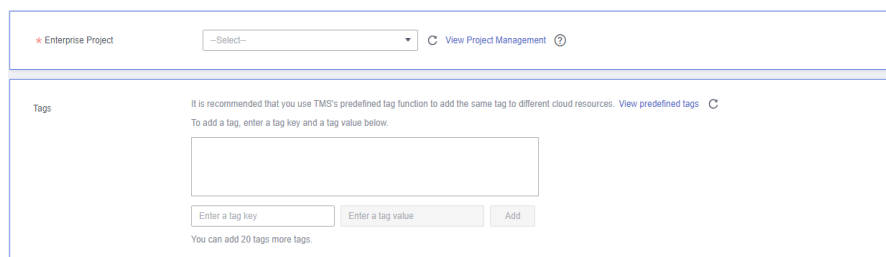


Table 3-269 Enterprise Project and Tags

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

 **NOTE**

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

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify 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 3-254 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 3-270 Source database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the source TiDB database. |
| Port | The port number must be an integer ranging from 1 to 65535, and the default value is 4000. |
| Database Username | The username for accessing the source TiDB database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

Figure 3-255 Kafka information

Kafka Information

IP Address ?

Kafka addresses must be IP addresses. They cannot be domain names or host names. In addition, listeners of all Kafka nodes must be configured in the IP address mode.

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

Table 3-271 Kafka information

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | IP address or domain name of Kafka DRS obtains the binlog generated by TiDB from Kafka and synchronizes the incremental data to the destination database in real time. |

Figure 3-256 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 3-272 Destination database settings

| Parameter | Description |
|-------------------|--|
| DB Instance Name | The GaussDB(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 3-257 Synchronization Mode

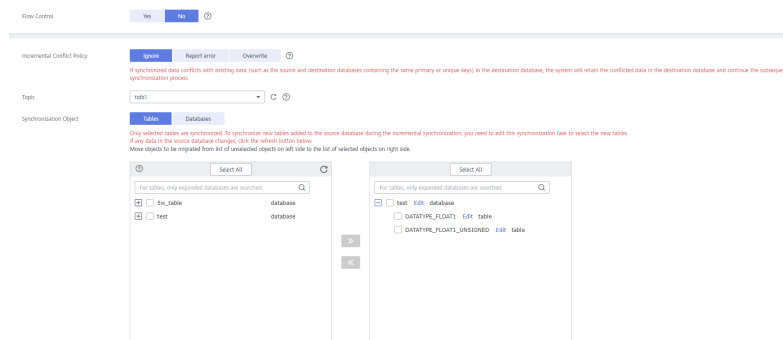
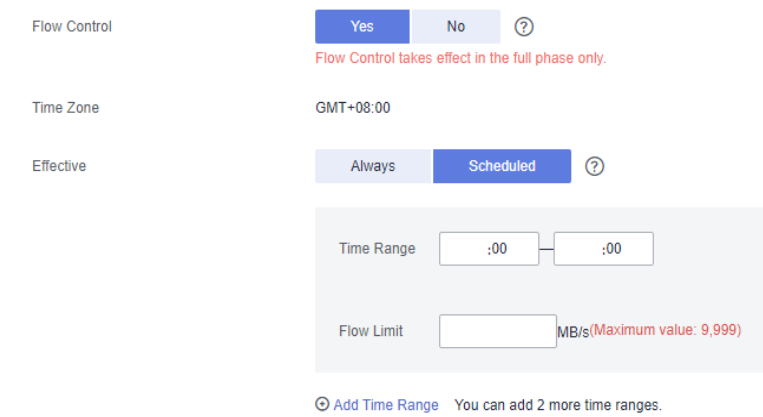



Table 3-273 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-258 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

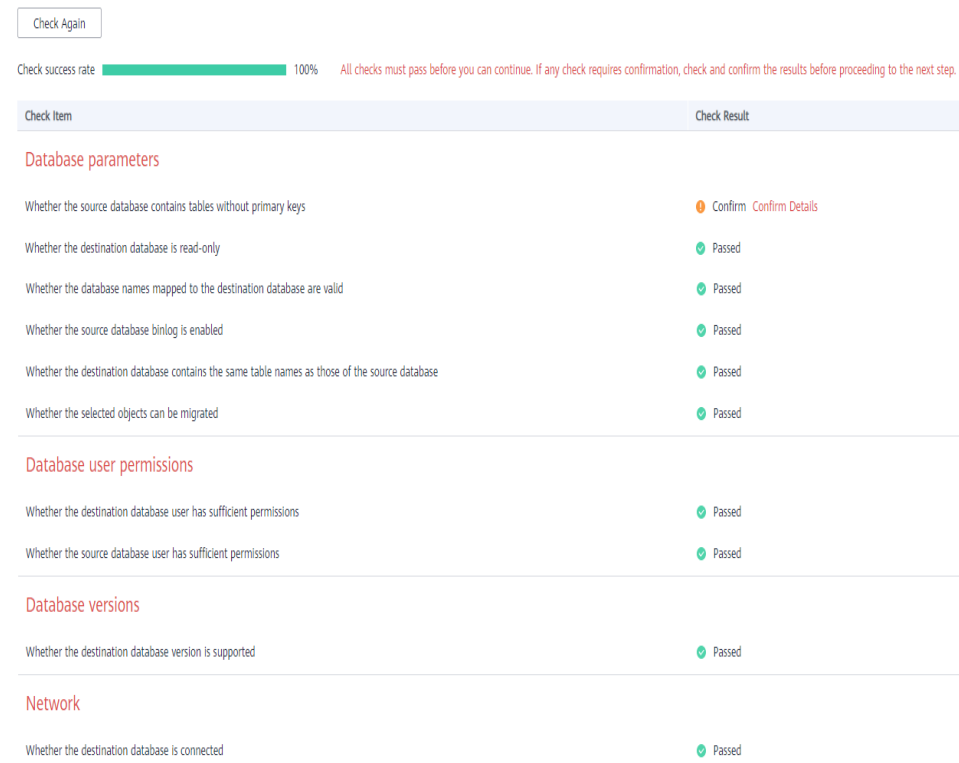
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-259 Pre-check



NOTE

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

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

Figure 3-260 Task startup settings

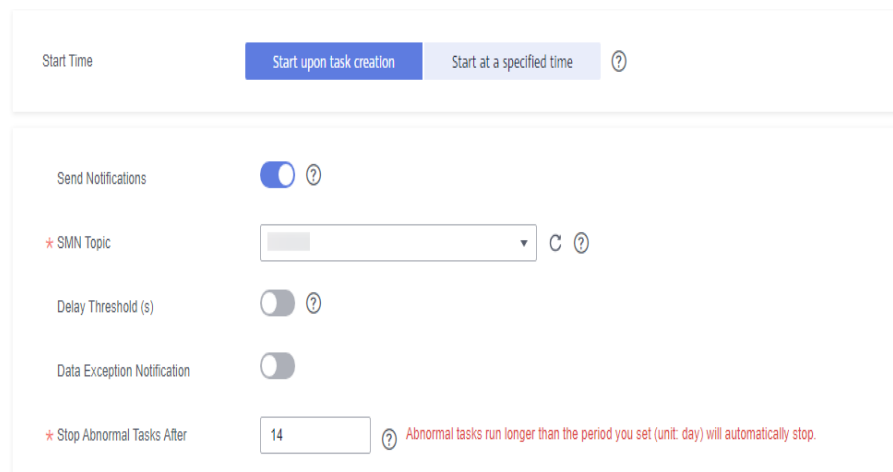



Table 3-274 Task startup settings

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

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

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

----End

3.26 From Microsoft SQL Server to GaussDB(DWS)

Supported Source and Destination Databases

Table 3-275 Supported databases

| Source DB | Destination DB |
|---|--|
| <ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) | <ul style="list-style-type: none"> • GaussDB(DWS) cluster |

Supported Synchronization Objects

Table 3-276 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-276 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 column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. - The number of tables selected for the synchronization object cannot exceed 1000. If there are more than 1000 tables, you are advised to synchronize them in batches. (Create a new task after the synchronization task is complete.) - Transparent Data Encryption (TDE) encrypted databases in the source instance cannot be synchronized. If you do not need to synchronize TDE encrypted databases, deselect them. If you need to synchronize TDE encrypted databases, disable TDE first. - Column encryption is not supported. - Auto-increment columns cannot be synchronized. ● 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 3-277](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-277 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

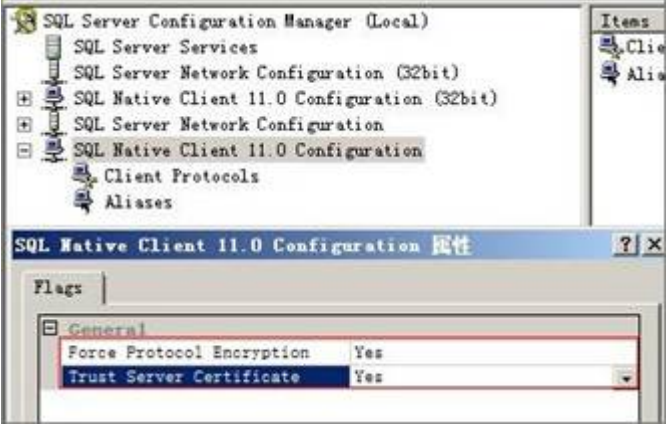
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- 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.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 3-278 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 3-261. <p style="text-align: center;">Figure 3-261 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database 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 |

| 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"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. |
| 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. |

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 3-262 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 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." The Task Name and Description fields have help icons. At the bottom right of the form, there is a character count "0/256".

Table 3-279 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

| 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 3-263 Synchronization instance details

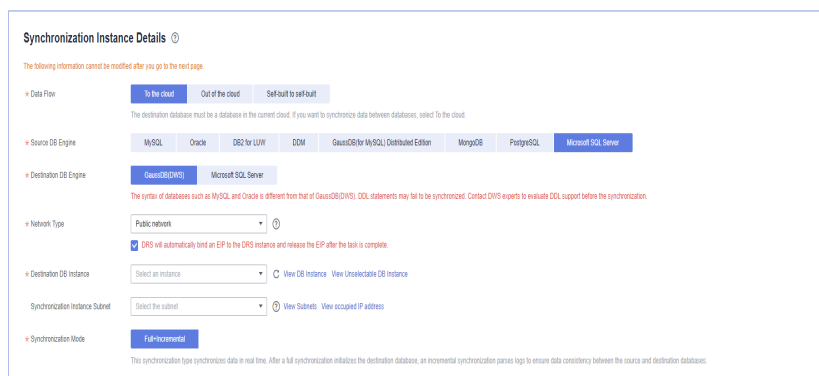


Table 3-280 Synchronization instance settings

| 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. <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| 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 Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. |

- AZ

Figure 3-264 AZ



Table 3-281 Task AZ

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

- Enterprise Project and Tags

Figure 3-265 Enterprise projects and tags

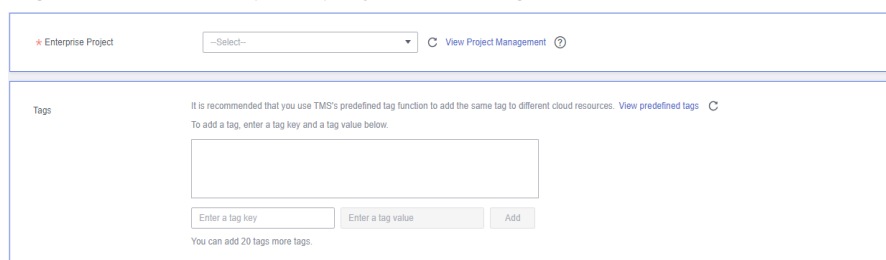


Table 3-282 Enterprise Project and Tags

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

 **NOTE**

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

Step 3 On the **Configure Source and Destination Databases** page, 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 3-266 Source database information

Source Database

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

Database Type Self-built on ECS **RDS DB Instance**

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

Database Username

Database Password

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

Table 3-283 Source database settings

| Parameter | Description |
|-------------------|---|
| Database Type | Select RDS DB instance . |
| DB Instance Name | Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance. |
| Database Username | The username for logging in to the source Microsoft SQL Server database. |
| Database Password | The password for the database username. |

Figure 3-267 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

● Test successful

Table 3-284 Destination database settings

| Parameter | Description |
|-------------------|---|
| DB Instance Name | The GaussDB(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 3-268 Synchronization mode

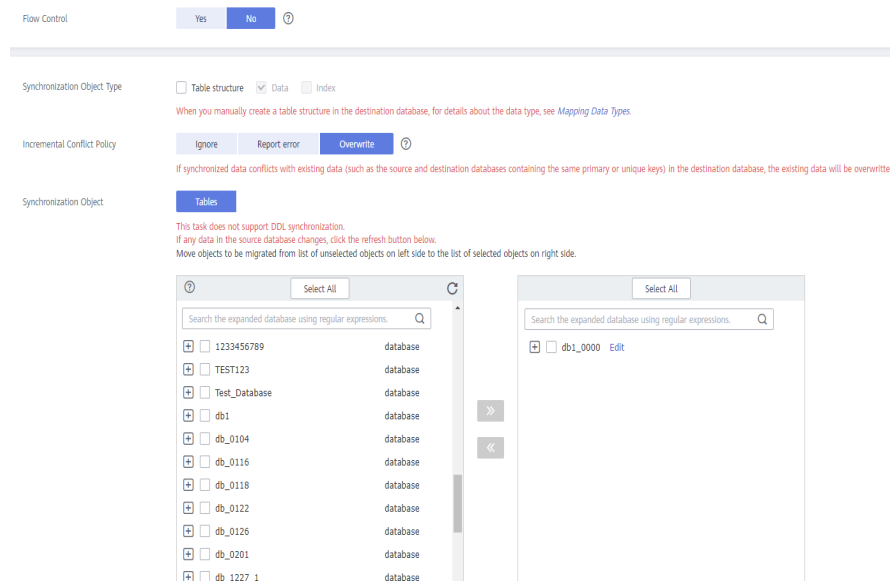
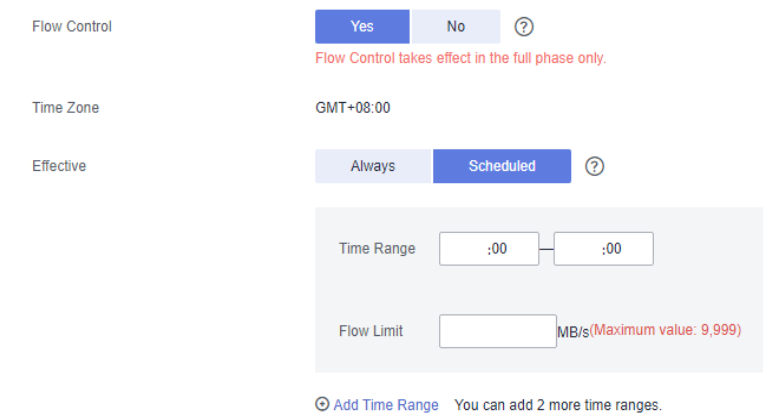



Table 3-285 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-269 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

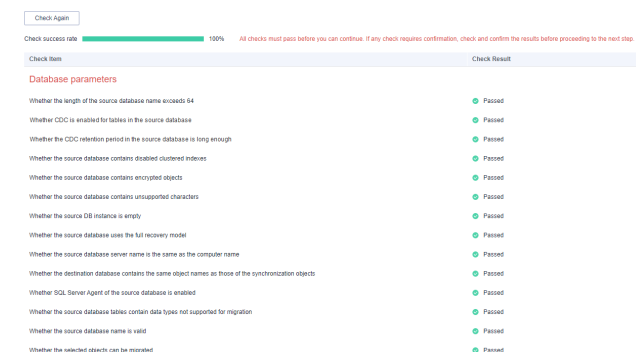
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-270 Pre-check



NOTE

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

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

Figure 3-271 Task startup settings

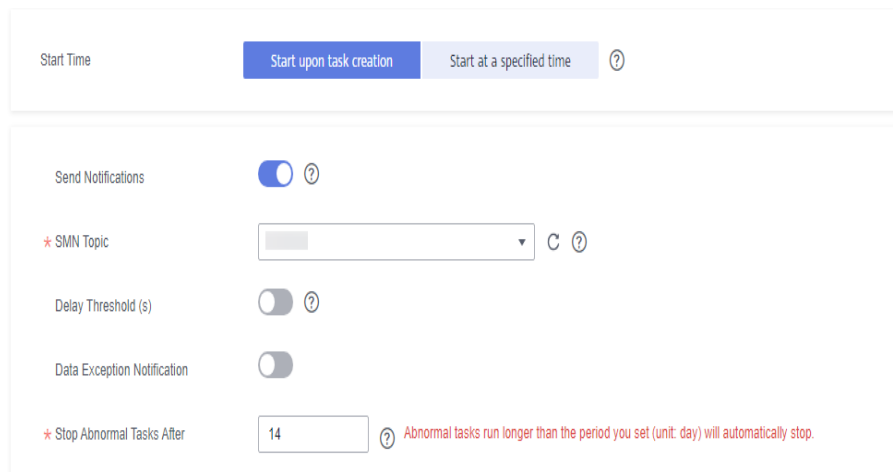



Table 3-286 Task startup settings

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

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

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

----End

3.27 From Microsoft SQL Server to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 3-287 Supported databases

| Source DB | Destination DB |
|---|-------------------------|
| <ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) | GaussDB primary/standby |

Supported Synchronization Objects

[Table 3-288](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-288 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. - Auto-increment columns cannot be synchronized. ● 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 3-289](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

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

Table 3-289 Database account permission

| Type | Full | Incremental and Full+Incremental |
|---------------------------|---|--|
| Source database user | Sysadmin permission; DB_DATAREADER or DB_OWNER permission for a database to be synchronized | Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized |
| Destination database user | <ul style="list-style-type: none"> • Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. • Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. • Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a

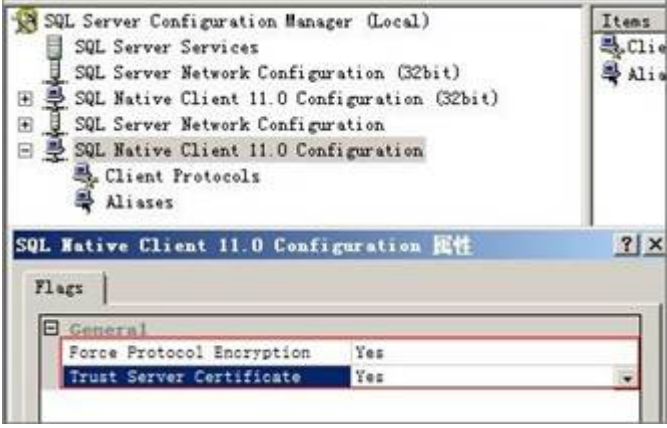
synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 3-290 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 3-272. <p style="text-align: center;">Figure 3-272 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database 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. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is |

| Type | Constraints |
|-----------------------------|--|
| | <p>required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later.</p> <ul style="list-style-type: none"> - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. - If the destination database is in Oracle compatibility mode, empty strings written to the destination database are processed as NULL. If the source database contains empty strings and has the NOT NULL constraint, data fails to be written to the destination database. |
| 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. • Tables whose primary key type is blob, text, clob, nlob, or bytea cannot be deleted or updated. • When adding synchronization objects in an incremental task, you are not advised to perform DML operations on the newly added synchronization tables before the subtask change is complete. Otherwise, some data may fail to be synchronized to the destination database. |

| Type | Constraints |
|----------------------------|---|
| Synchronization comparison | <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. |

Procedure

This section uses Microsoft SQL Server to 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 3-273 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]

Table 3-291 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

| 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 3-274 Synchronization instance details

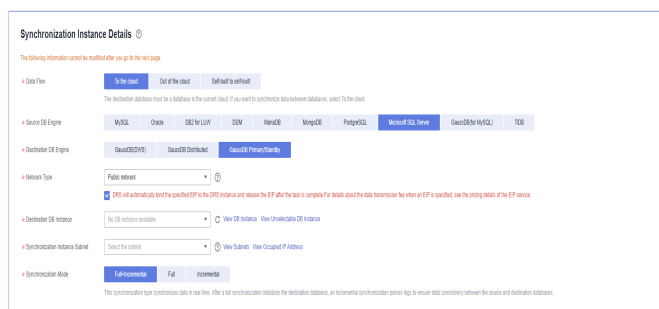


Table 3-292 Synchronization instance settings

| Parameter | Description |
|-------------------------|---|
| Data Flow | Select To the cloud . |
| Source DB Engine | Select Microsoft SQL Server . |
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example. <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | An available GaussDB primary/standby instance. |

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

- AZ

Figure 3-275 AZ



Table 3-293 Task AZ

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

- Enterprise Project and Tags

Figure 3-276 Enterprise projects and tags

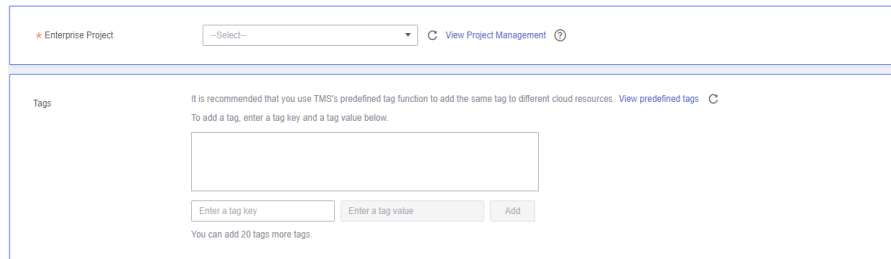


Table 3-294 Enterprise Project and Tags

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

NOTE

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

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

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

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

Figure 3-277 Source database information

Source Database

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

Database Type Self-built on ECS **RDS DB Instance**

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

Database Username

Database Password

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

Table 3-295 Source database settings

| Parameter | Description |
|-------------------|---|
| Database Type | Select RDS DB instance . |
| DB Instance Name | Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance. |
| Database Username | The username for logging in to the source Microsoft SQL Server database. |
| Database Password | The password for the database username. |

Figure 3-278 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Test Connection ✔ Test successful

Table 3-296 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 3-279 Synchronization mode

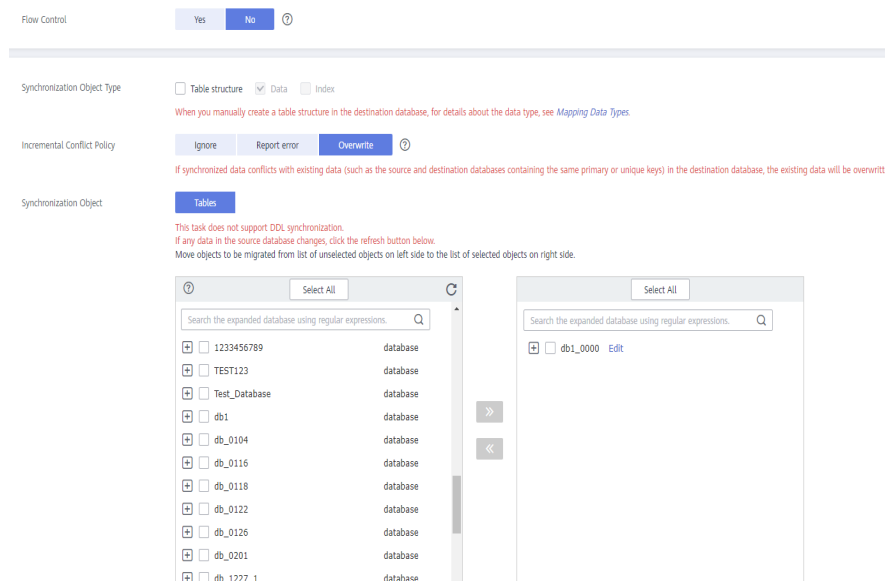
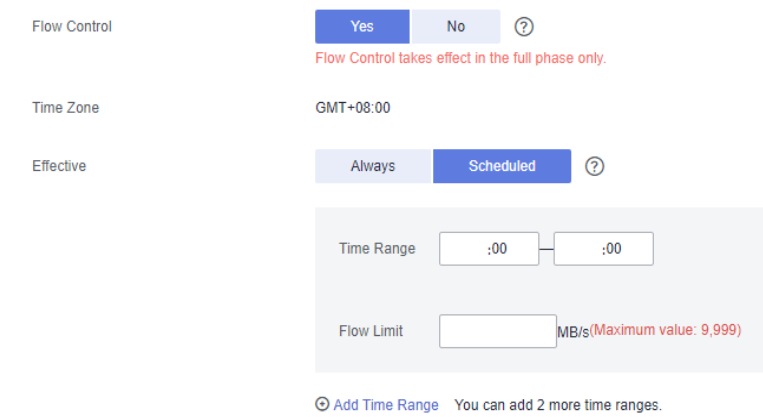



Table 3-297 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-280 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

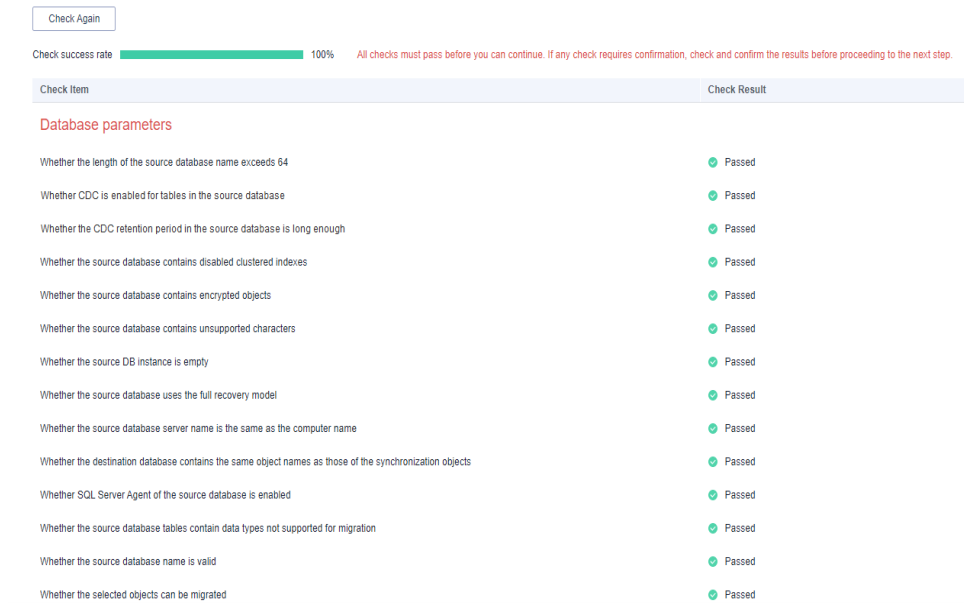
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-281 Pre-check



NOTE

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

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

Figure 3-282 Task startup settings

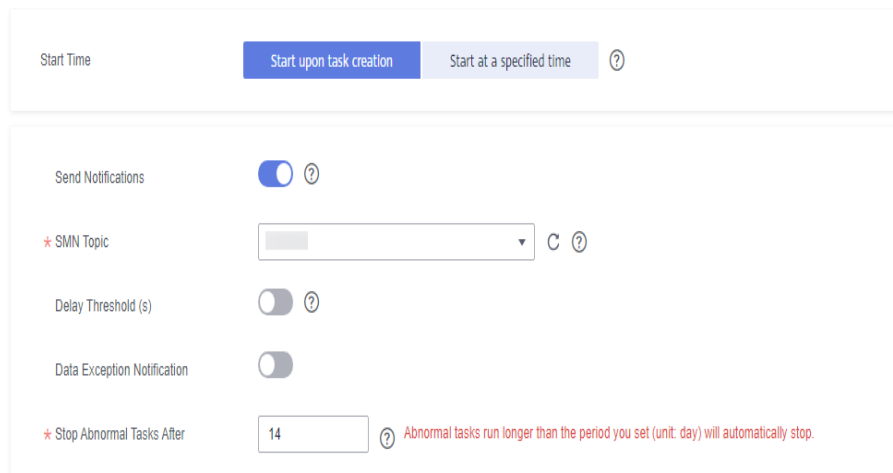



Table 3-298 Task startup settings

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

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

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

----End

3.28 From Microsoft SQL Server to GaussDB Distributed

Supported Source and Destination Databases

Table 3-299 Supported databases

| Source DB | Destination DB |
|---|---------------------|
| <ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) | GaussDB distributed |

Supported Synchronization Objects

Table 3-300 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-300 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. - Auto-increment columns cannot be synchronized. ● 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 3-301](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

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

Table 3-301 Database account permission

| Type | Full | Full+Incremental |
|---------------------------|---|--|
| Source database user | Sysadmin permission; DB_DATAREADER or DB_OWNER permission for a database to be synchronized | Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized |
| Destination database user | <ul style="list-style-type: none"> • Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. • Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. • Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a

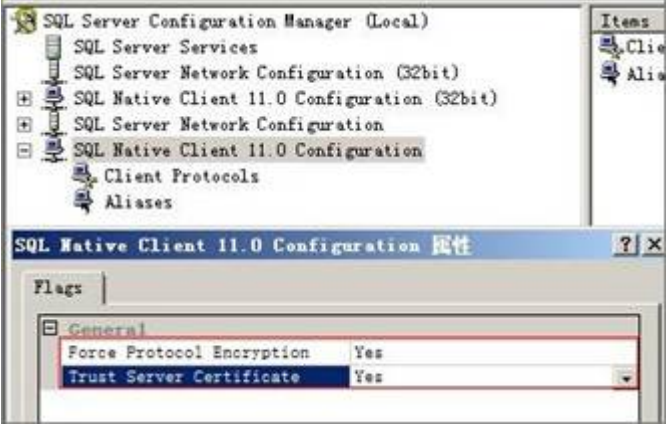
synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 3-302 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 3-283. <p style="text-align: center;">Figure 3-283 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database 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. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is |

| Type | Constraints |
|-----------------------------|--|
| | <p>required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later.</p> <ul style="list-style-type: none"> - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. - If the destination database is in Oracle compatibility mode, empty strings written to the destination database are processed as NULL. If the source database contains empty strings and has the NOT NULL constraint, data fails to be written to the destination database. |
| 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. • Tables whose primary key type is blob, text, clob, nlob, or bytea cannot be deleted or updated. • When adding synchronization objects in an incremental task, you are not advised to perform DML operations on the newly added synchronization tables before the subtask change is complete. Otherwise, some data may fail to be synchronized to the destination database. |

| Type | Constraints |
|----------------------------|---|
| Synchronization comparison | <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. |

Procedure

This section uses Microsoft SQL Server to 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 3-284 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is an orange warning banner 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 includes:

- Region:** A dropdown menu with a selected option.
- Project:** A dropdown menu with a selected option.
- * 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 3-303 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

| 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 3-285 Synchronization instance details

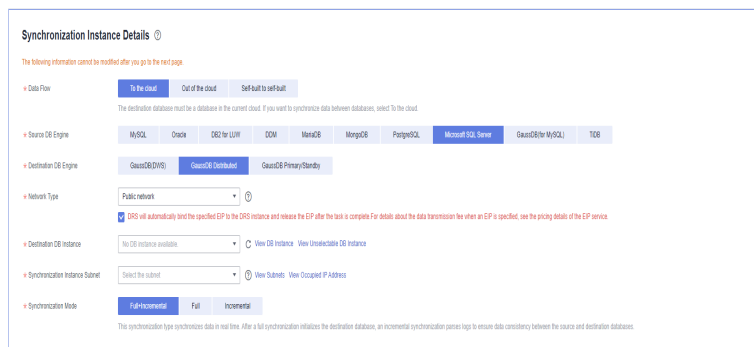


Table 3-304 Synchronization instance settings

| 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. <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | An available GaussDB distributed instance. |

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

- AZ

Figure 3-286 AZ




Table 3-305 Task AZ

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

- Enterprise Project and Tags

Figure 3-287 Enterprise projects and tags




Table 3-306 Enterprise Project and Tags

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

NOTE

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

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

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

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

Figure 3-288 Source database information

Source Database

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

Database Type Self-built on ECS **RDS DB Instance**

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

Database Username

Database Password

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

Table 3-307 Source database settings

| Parameter | Description |
|-------------------|---|
| Database Type | Select RDS DB instance . |
| DB Instance Name | Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance. |
| Database Username | The username for logging in to the source Microsoft SQL Server database. |
| Database Password | The password for the database username. |

Figure 3-289 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Test Connection ✔ Test successful

Table 3-308 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 3-290 Synchronization mode

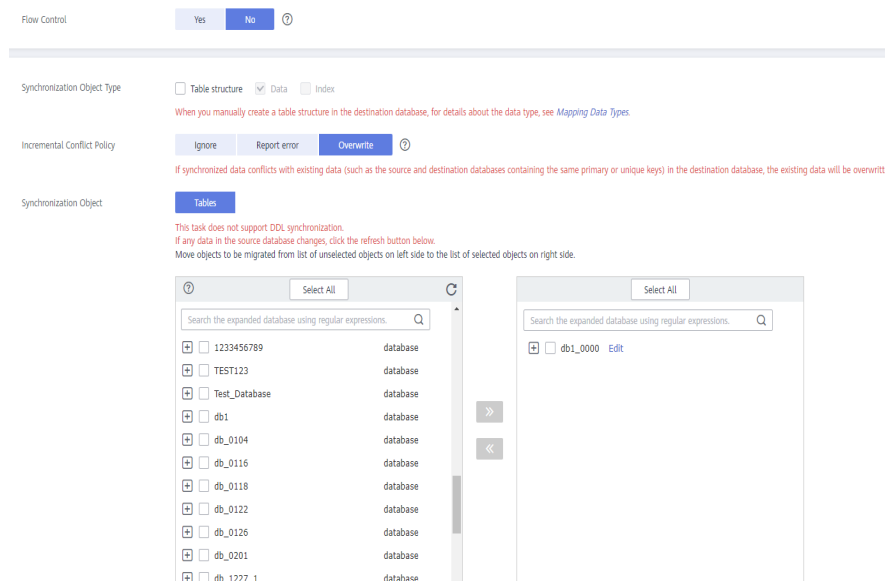
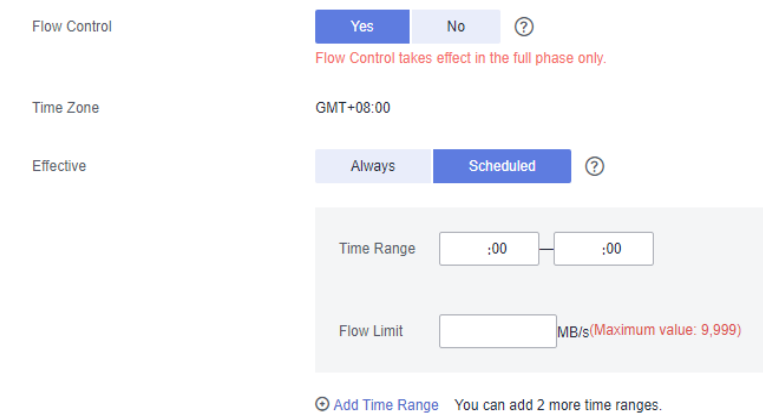



Table 3-309 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-291 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

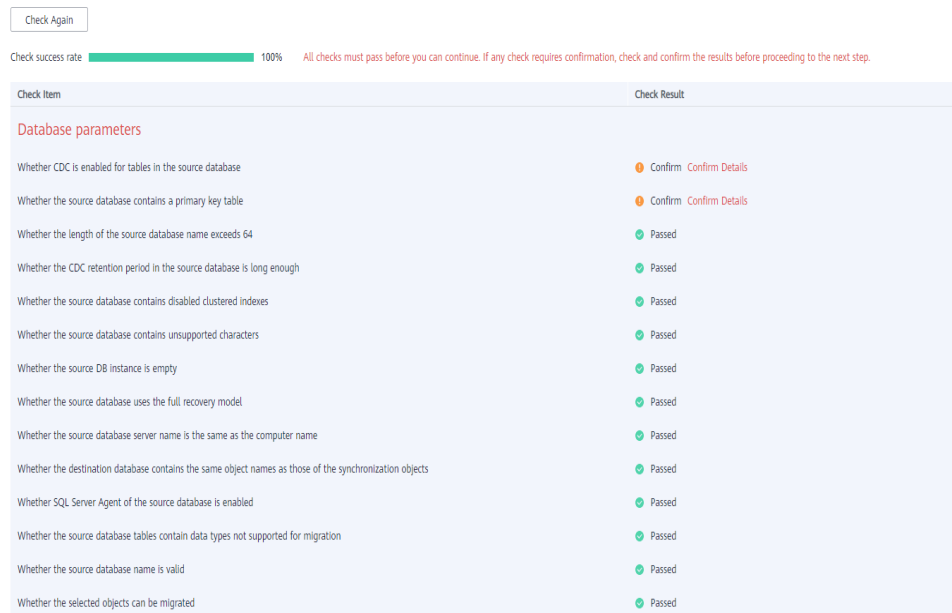
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-292 Pre-check



NOTE

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

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

Figure 3-293 Task startup settings

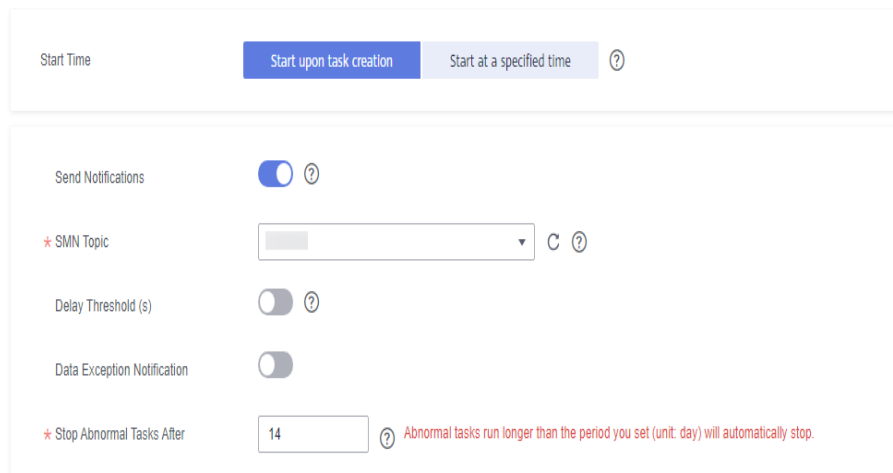



Table 3-310 Task startup settings

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

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

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

----End

3.29 From Microsoft SQL Server to Microsoft SQL Server

Supported Source and Destination Databases

Table 3-311 Supported databases

| Source DB | Destination DB |
|---|--|
| <ul style="list-style-type: none"> • On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • Microsoft SQL Server on other clouds (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) | <ul style="list-style-type: none"> • RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) <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 3-312](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-312 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 selected tables 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. - Auto-increment columns cannot be synchronized. • 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 3-313](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-313 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 | <p>If the destination end does not contain databases, the destination database user must have the create any database permission.</p> <p>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.</p> |

Suggestions

 CAUTION

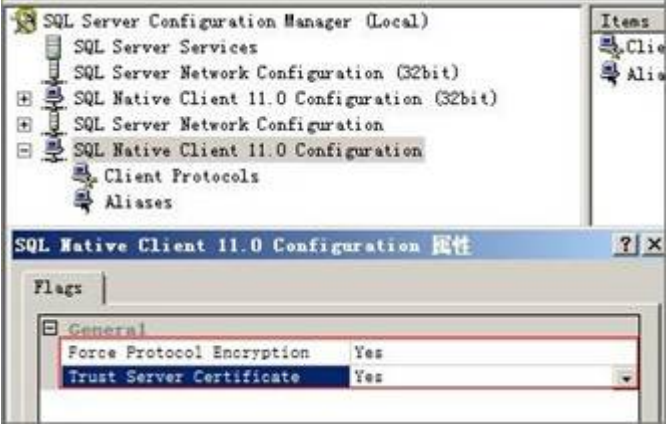
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - 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 3-314 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 3-294. <p style="text-align: center;">Figure 3-294 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● 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. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. - Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server |

| Type | Constraints |
|-----------------------------|---|
| | <p>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. |

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 3-295 Synchronization task information

Table 3-315 Task information

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

- Synchronization instance details

Figure 3-296 Synchronization instance details

Table 3-316 Synchronization instance information

| Parameter | Description |
|------------------|--------------------------------------|
| Data Flow | Select To the cloud . |
| Source DB Engine | Select Microsoft SQL Server . |

| Parameter | Description |
|---------------------------------|---|
| 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. <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| 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 Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Synchronization Mode | The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . <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. |

- AZ

Figure 3-297 AZ



Table 3-317 Task AZ

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

- Enterprise Project and Tags

Figure 3-298 Enterprise projects and tags

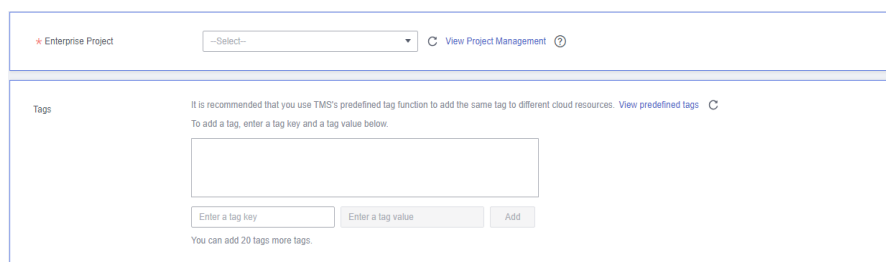


Table 3-318 Enterprise Project and Tags

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

NOTE

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

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

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

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

Figure 3-299 Source database information

Source Database

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

Database Type: Self-built on ECS RDS DB Instance

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

Database Username:

Database Password:

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

Table 3-319 Source database settings

| Parameter | Description |
|-------------------|---|
| Database Type | Select RDS DB instance . |
| DB Instance Name | Select the Microsoft SQL Server DB instance to be synchronized as the source DB instance. |
| Database Username | The username for logging in to the source Microsoft SQL Server database. |
| Database Password | The password for the database username. |

Figure 3-300 Destination database information

Table 3-320 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. |

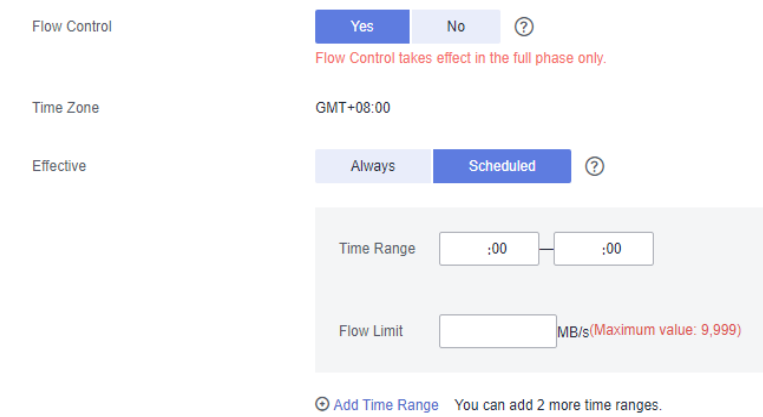
NOTE


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 3-301 Synchronization Mode

Table 3-321 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-302 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Common index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is mandatory. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

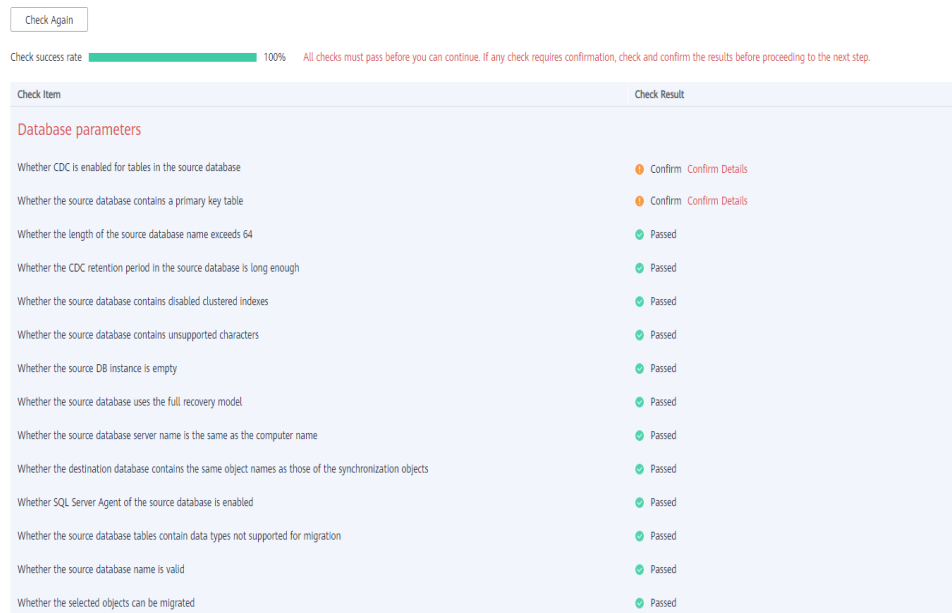
Step 5 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-303 Pre-check



NOTE

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

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

Figure 3-304 Task startup settings





Table 3-322 Task startup settings

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

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

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

----End

3.30 From MongoDB to DDS

Supported Source and Destination Databases

Table 3-323 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none"> • On-premises MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4) • ECS-hosted MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4) • Other cloud MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4) • DDS DB instances (versions 3.2, 3.4, 4.0, 4.2, and 4.4) | <ul style="list-style-type: none"> • DDS DB instances (versions 3.4, 4.0, 4.2, and 4.4) <p>NOTE The destination database version must be the same as or later than the source database version.</p> |

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.
- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full synchronization, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-324 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> |

| Type | Restrictions |
|------------------------------------|--|
| 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. ● If the source is a cluster instance of version 4.4, composite hash indexes and composite hash shard keys are not supported. ● If the source is a replica set instance of version 4.4, composite hash indexes are not supported. ● The mapped synchronization object name can contain 1 to 63 characters. The following characters are not allowed <code>\. "\$ <></code> ● Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the synchronization. |
| Source database | <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 or view name cannot start with system. or contain the dollar sign (<code>\$</code>). ● The source cannot be a GeminiDB 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 3-305 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 3-325 Task information

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

- Synchronization instance information

Figure 3-306 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 3-326 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 , Public network and VPN or Direct Connect . Public network is used as an example. <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Destination DB Instance | Select the DB instance you have created. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |

| Parameter | Description |
|-------------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <p>Full+Incremental: This synchronization type allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases.</p> <p>NOTE 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 Type | <p>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.</p> |

- Task type

Figure 3-307 Task type



Table 3-327 Task type information

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

- Enterprise Project and Tags

Figure 3-308 Enterprise projects and tags




Table 3-328 Enterprise Project and Tags

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

NOTE

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

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

- Source database information

Figure 3-309 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 3-329 Source database settings

| Parameter | Description |
|---------------------------|---|
| Database Types | The source database type can be Self-built on ECS or a DDS DB Instance . In this example, DDS DB Instance is selected. |
| VPC | This option is available if the source database type is Self-built on ECS . VPC is used to isolate networks for different services. |
| Subnet | This option is available if the source database type is Self-built on ECS . Subnet provides dedicated and isolated networks for different services and controls access. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet. |
| IP Address or Domain Name | The IP address or domain name of the source database. This option is available if the source database type is Self-built on ECS . |
| DB Instance Name | The DDS DB instance as the source. This option is available if the source database type is DDS DB Instance . |
| Authentication Database | The name of the authentication database. For example: The default authentication database of DDS instance is admin . |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

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

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-310 Destination database information

Destination Database

DB Instance Name

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 3-330 Destination database settings

| Parameter | Description |
|-------------------------|---|
| DB Instance Name | The default value is the DB instance selected when the synchronization task is created and cannot be changed. |
| Authentication Database | The name of the authentication database. For example: The default authentication database of DDS instance is admin . |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

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

 **NOTE**

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

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

Figure 3-311 Synchronization objects

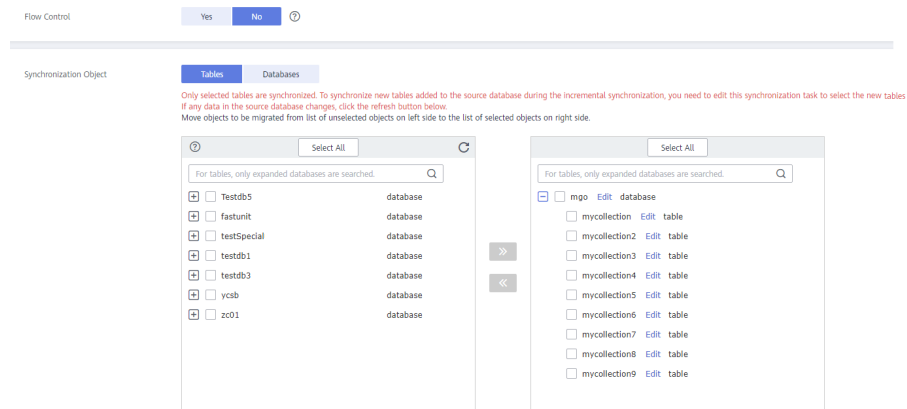
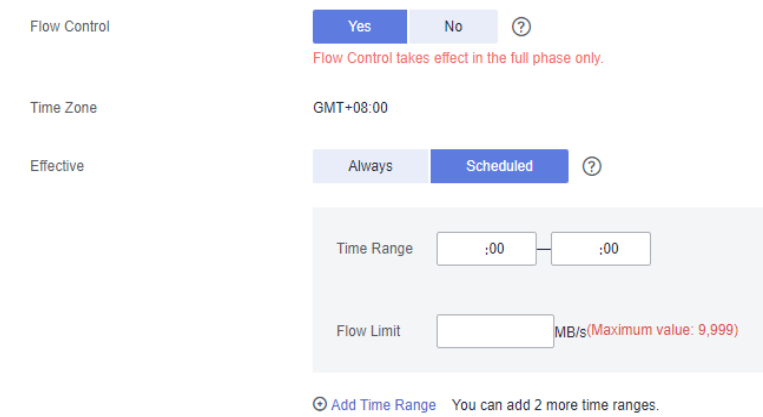




Table 3-331 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-312 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|------------------------|--|
| Synchronization Object | <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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

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

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

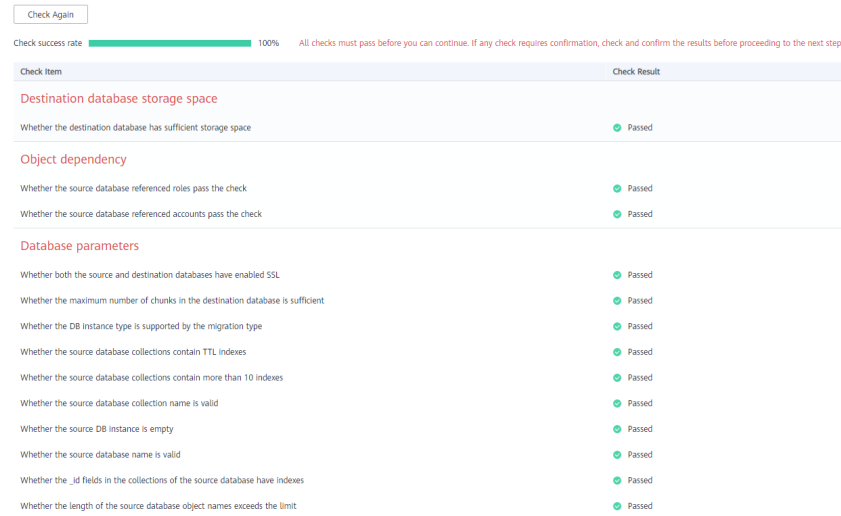
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Figure 3-313 Pre-check



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

Figure 3-314 Task startup settings

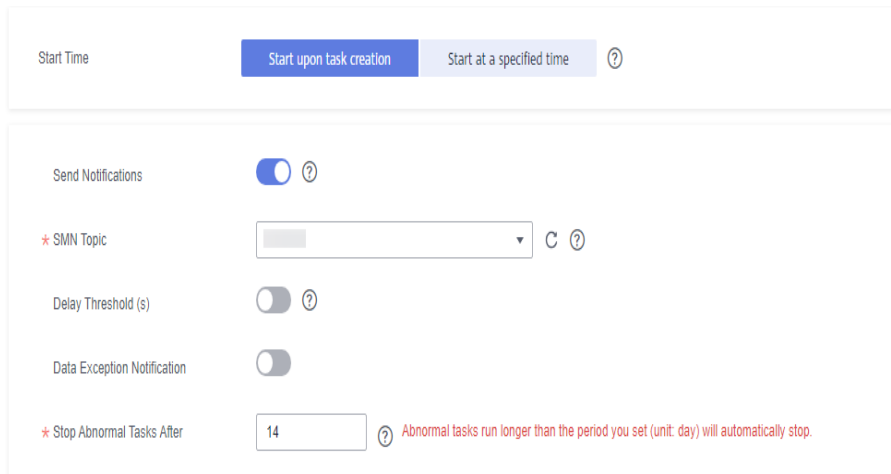



Table 3-332 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |

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

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

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

----End

3.31 From MariaDB to MariaDB

Supported Source and Destination Databases

Table 3-333 Supported databases

| Source DB | Destination DB |
|---|----------------------|
| <ul style="list-style-type: none">• On-premises MariaDB 10.3, 10.4, and 10.5• ECS-hosted MariaDB 10.3, 10.4, and 10.5• Other cloud MariaDB 10.3, 10.4, and 10.5 | RDS for MariaDB 10.5 |

Supported Synchronization Objects

Table 3-334 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-334 Supported synchronization objects

| Type | Precautions |
|---------|--|
| Objects | <ul style="list-style-type: none"> ● Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● Events and triggers cannot be synchronized. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● 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 (<code>_</code>), or the mapped database table name contains hyphens (<code>-</code>) and number signs (<code>#</code>), the name length cannot exceed 42 characters. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur. |

Database Account Permission Requirements

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

tasks require different permissions. For details, see [Table 3-335](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

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

Table 3-335 Database account permission

| Type | Full+Incremental |
|---------------------------|---|
| Source database user | SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions) |
| Destination database user | The root account of RDS for MariaDB has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 3-336 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - During the incremental synchronization, the binlog of the source MariaDB database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. - If the expire_logs_days value of the source database is set to 0, the synchronization may fail. - Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. - During an incremental synchronization, the source MariaDB database server_id must be set to a value ranging from 1 to 4294967296. ● Source database object requirements: <ul style="list-style-type: none"> - The source database names cannot contain non-ASCII characters, or the following characters: '<>/\'" - 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. ● Destination database parameter requirements: <ul style="list-style-type: none"> - Data cannot be synchronized from a newer version database to an older version database. - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. - If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails. ● Destination database object requirements: |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> - The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal. - The destination DB instance must have sufficient storage space. - If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. - The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') <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. - 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 DB instances are RDS for MariaDB instances, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. - If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection. - If the source database version is earlier than 10.4.3 and database-level synchronization is selected, after the task is started, do not use table creation statements (for example, create table t1 (c1 json)) containing JSON data type to create table structures or use column adding statements (for example, alter table t1 add column c1 json) containing JSON data type in the database selected for the synchronization object. Otherwise, data will be inconsistent. - Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. - The destination table can contain more columns than the source table. However, the following failures must be avoided: |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. - The source database does not support point-in-time recovery (PITR). - The destination database cannot be restored to a point in time when a full synchronization was being performed. - Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. - Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. - The partitioned table does not support column mapping. - Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - After a task is created, the destination database cannot be set to read-only. - If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |

| Type | Restrictions |
|-----------------------------|---|
| Full synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization. ● During full data synchronization, a lot of binlogs are generated in the destination database, occupying too much storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. |
| 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 statements can be synchronized during incremental synchronization. ● Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● You can add additional objects during an incremental synchronization. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

Procedure

This section describes how to use DRS to configure a MariaDB 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 3-315 Synchronization task information

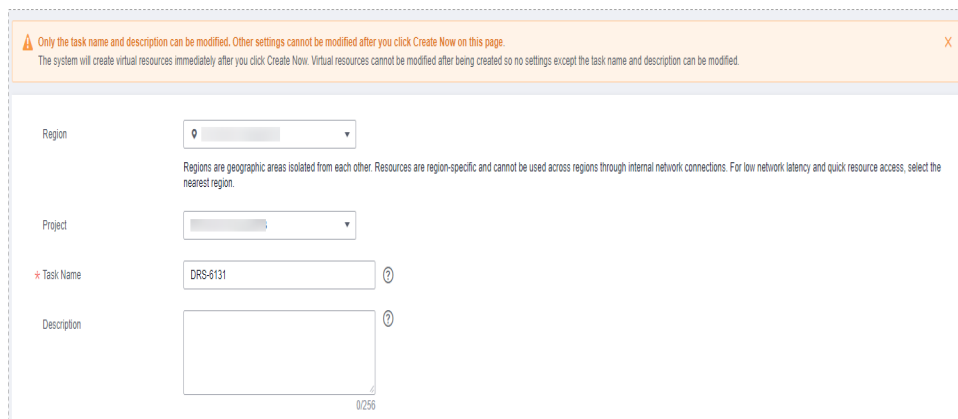


Table 3-337 Task information

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

- Synchronization instance information

Figure 3-316 Synchronization instance information

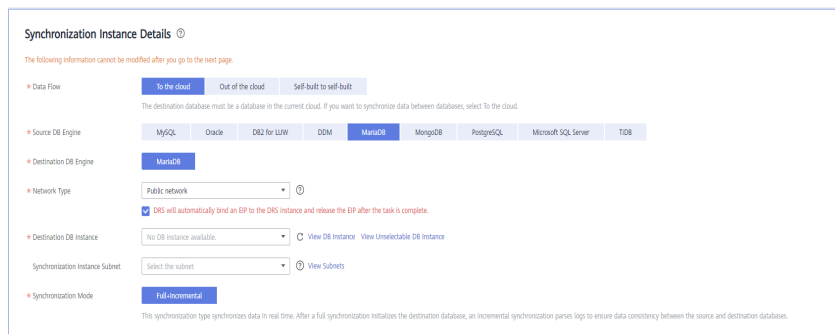


Table 3-338 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <p>Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases.</p> <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |

- Task type

Figure 3-317 Task type



Table 3-339 Task type information

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

- Enterprise Project and Tags

Figure 3-318 Enterprise projects and tags

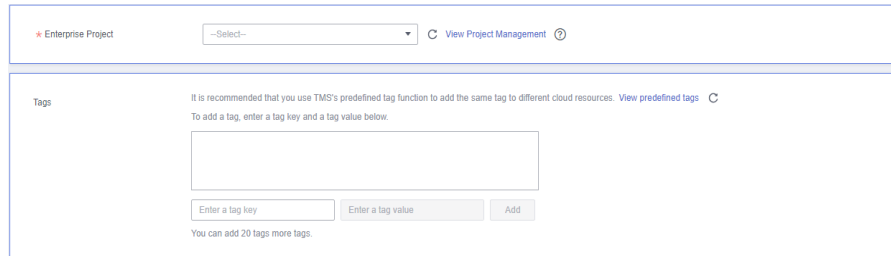


Table 3-340 Enterprise Project and Tags

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

NOTE

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

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

- Source database information

Figure 3-319 Source database information

Table 3-341 Source database information

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk. |

NOTE

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-320 Destination database information

Destination Database

DB Instance Name ddt-shard-syn-ia

Database Username

Database Password

Table 3-342 Destination database settings

| Parameter | Description |
|-------------------|--|
| DB Instance Name | The RDS for MariaDB instance you selected when creating the task. The parameter cannot be changed. |
| Database Username | The username for accessing the destination database. |
| Database Password | Password for logging in to the destination database. |

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 3-321 Synchronization objects

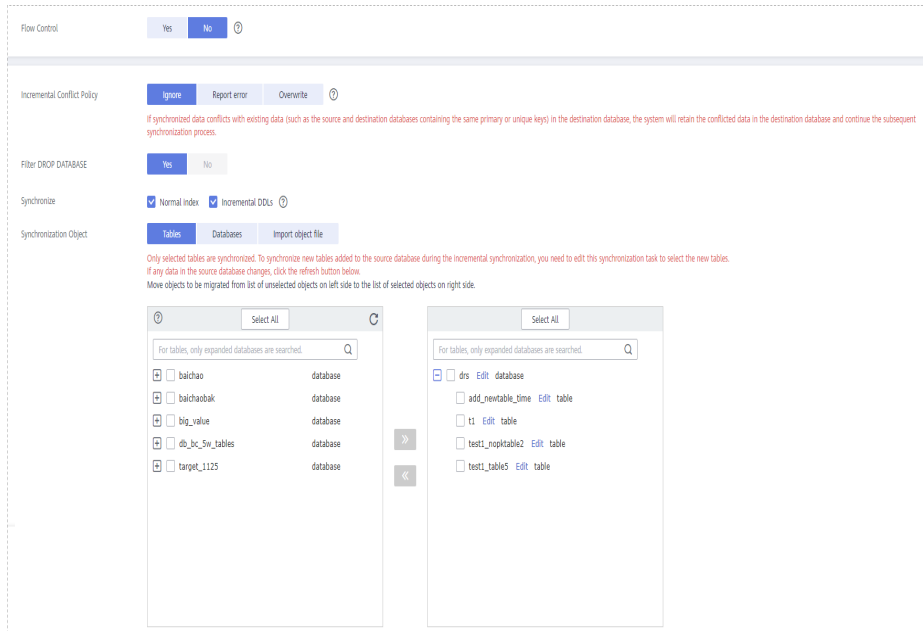
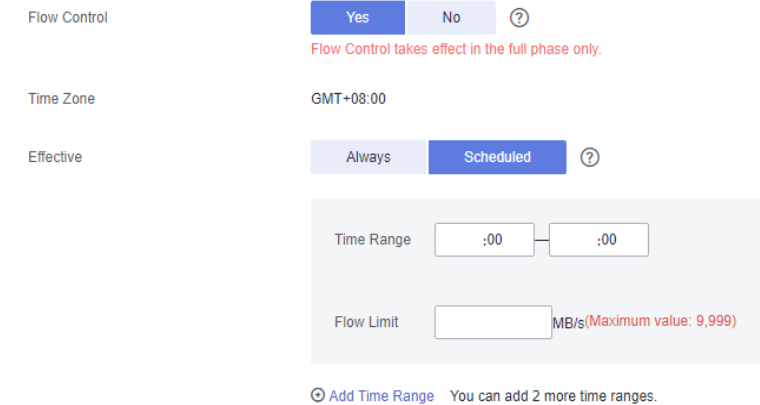



Table 3-343 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-322 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize data based on service requirements.</p> |

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

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

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

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

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

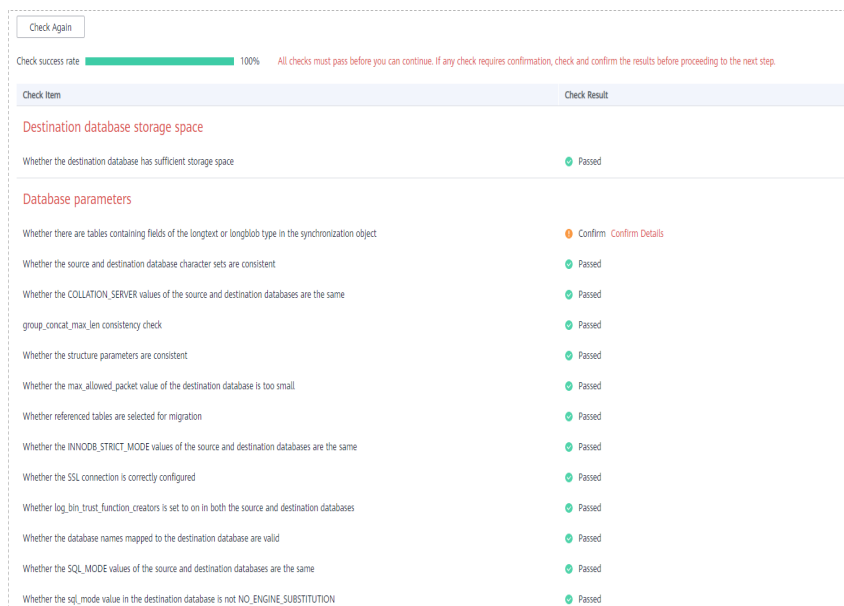
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Figure 3-323 Pre-check



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

Figure 3-324 Task startup settings

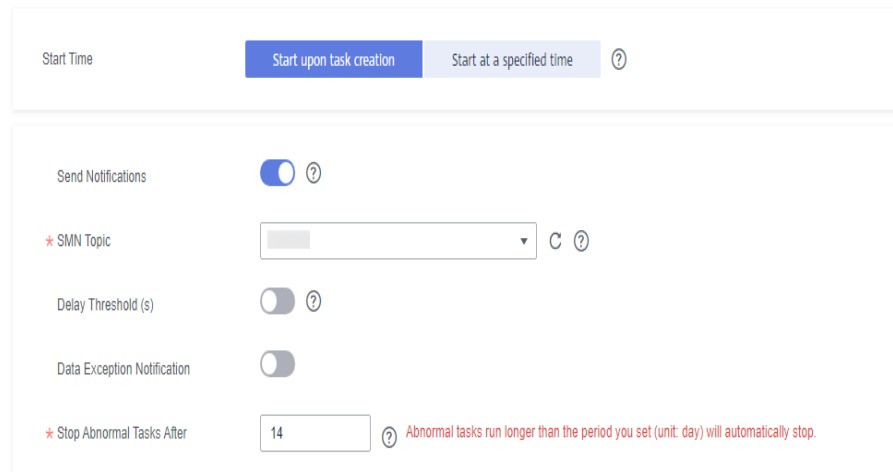



Table 3-344 Task startup settings

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

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

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

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

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

3.32 From MariaDB to MySQL

Supported Source and Destination Databases

Table 3-345 Supported databases

| Source DB | Destination DB |
|--|--|
| <ul style="list-style-type: none"> On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 Other cloud MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 | RDS for MySQL 5.5, 5.6, 5.7, and 8.0 NOTE <ul style="list-style-type: none"> If the source database version is MariaDB 10.0 or 10.1, you are advised to select MySQL 5.6 or later as the destination database. If the source database version is MariaDB 10.2, 10.3 or 10.4, you are advised to select MySQL 5.7 or later as the destination database. If the source database version is MariaDB 10.5, you are advised to select MySQL 8.0 as the destination database. |

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

[Table 3-346](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-346 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none">• Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized.• Only MyISAM and InnoDB tables can be synchronized.• Events and triggers cannot be synchronized.• Table names cannot be mapped for tables on which views, stored procedures, and functions depend.• When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized.• 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.• The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent.<ul style="list-style-type: none">- If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out.- If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out.- You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent.• If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur. |

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-347](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

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

Table 3-347 Database account permission

| Type | Incremental and Full+Incremental |
|---------------------------|--|
| Source database user | SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions) |
| Destination database user | The root account of 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. |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-Level Comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

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 3-348 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - During the incremental synchronization, the binlog of the source MariaDB database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is a self-managed MariaDB 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 MariaDB instance, set the binlog retention period by following the instructions provided in RDS User Guide. - Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. - During an incremental synchronization, the source MariaDB database server_id must be set to a value ranging from 1 to 4294967296. ● Source database object requirements: <ul style="list-style-type: none"> - The source database names cannot contain non-ASCII characters, or the following characters: '<>/\'' - 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. ● Destination database parameter requirements: <ul style="list-style-type: none"> - Data cannot be synchronized from a newer version database to an older version database. - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> - 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. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal. - The destination DB instance must have sufficient storage space. - If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. - The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see What Are Syntax Differences Between MySQL or MariaDB Versions?. - Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys - If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. - If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. - If the sources and destinations are RDS instances, database mapping is required. - The source and destination databases cannot contain tables that have the same names but do not have primary keys. |

| 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. - If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. - Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. - The source database does not support point-in-time recovery (PITR). - The destination database cannot be restored to a point in time when a full synchronization was being performed. - Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. - Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. - The partitioned table does not support column mapping. - After a task is created, the destination database cannot be set to read-only. - During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding |

| Type | Restrictions |
|----------------------|---|
| | <p>additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies.</p> <ul style="list-style-type: none"> - If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |
| Full synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • Data inconsistency may occur when the MyISAM table is modified during a full synchronization. • During full data synchronization, a lot of binlogs are generated in the destination database, occupying too much storage space. Therefore, during full data synchronization, only the latest five binlogs are retained in the destination database by default. After the full synchronization is complete, the retention period of binlogs in the destination database is restored to the value you configure. If you want to keep the binlog retention period of the destination database to be the value you specify due to service requirements, you need to submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. |

| 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 write data to the destination database. Otherwise, data may be inconsistent. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● You can add additional objects during an incremental synchronization. ● During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

Procedure

This section describes how to use DRS to configure a real-time synchronization task from MariaDB to RDS for MySQL over a public network.

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

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

- Task information description

Figure 3-325 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, 0/256]

Table 3-349 Task information

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

- Synchronization instance information

Figure 3-326 Synchronization instance information

Synchronization Instance Details

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

- Data Flow:** To the cloud, Out of the cloud, Self-suited self-suited
- Source DB Engine:** MySQL, Oracle, DB2 for LUW, CDCM, HuaweiDB, MongoDB, PostgreSQL, Microsoft SQL Server, GaussDB(for MySQL), TDB
- Destination DB Engine:** MySQL, MemDB, GaussDB(for MySQL)
- Network Type:** Public network
- DRS Task Type:** Single AZ, Dual AZ
- Destination DB Instance:** [Select an instance]
- Synchronization Instance Subnet:** [Select the subnet]
- IP Address Type:** IPv4, IPv4/IPv6 dual stack
- Synchronization Mode:** Full/Incremental, Incremental

Table 3-350 Synchronization instance settings

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

| Parameter | Description |
|---------------------------------|--|
| Destination DB Instance | The RDS for MySQL instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 3-327 Specifications





Table 3-351 Task type information

| Parameter | Description |
|-----------|---|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-328 AZ</p>  |

- Enterprise Project and Tags

Figure 3-329 Enterprise projects and tags




Table 3-352 Enterprise Project and Tags

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

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

 **NOTE**

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

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

- Source database information


Figure 3-330 Source database information

Source Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection


 Test successful

Table 3-353 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 |

| Parameter | Description |
|-------------------|---|
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk. |

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-331 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-354 Destination database settings

| Parameter | Description |
|------------------|--|
| DB Instance Name | The RDS for MySQL instance you selected when creating the task. The parameter cannot be changed. |

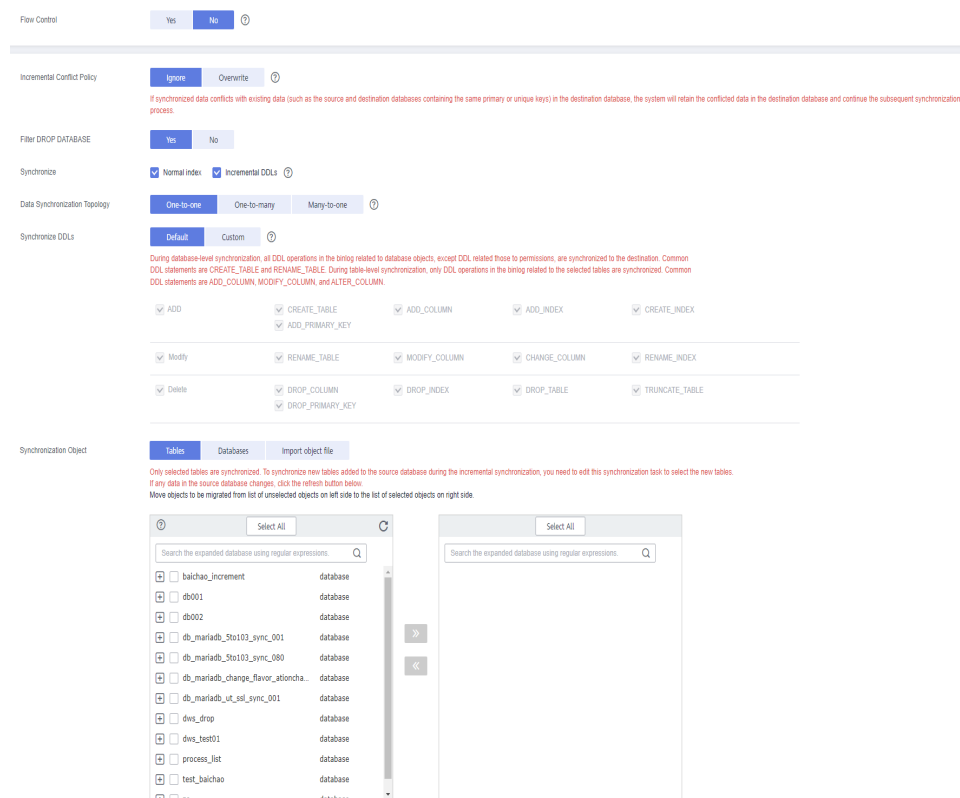
| Parameter | Description |
|-------------------|--|
| Database Username | The username for accessing the destination database. |
| Database Password | Password for logging in to the destination database. |

 **NOTE**

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

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

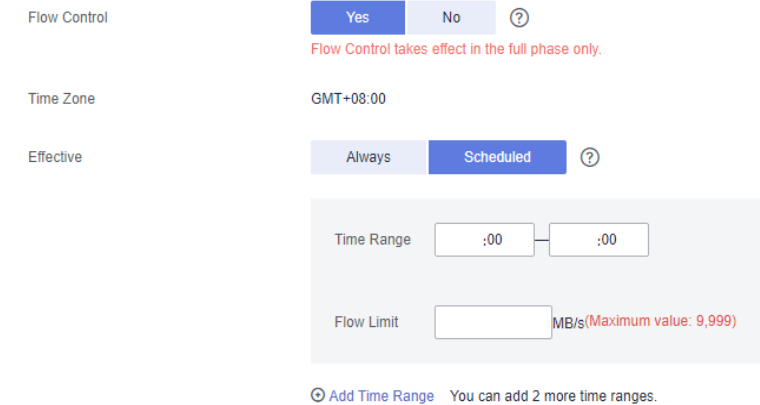
Figure 3-332 Synchronization objects



The screenshot displays the configuration interface for a synchronization task. Key sections include:


- Flow Control:** Buttons for 'Yes' and 'No'.
- Incremental Conflict Policy:** 'Ignore' and 'Overwrite' options.
- Filter DROP DATABASE:** 'Yes' and 'No' options.
- Synchronise:** 'Normal index' and 'Incremental DDLs' checkboxes.
- Data Synchronization Topology:** 'One-to-one', 'One-to-many', and 'Many-to-one' options.
- Synchronise DDLs:** 'Default' and 'Custom' options. A detailed list of DDL operations is shown with checkboxes, including ADD, Modify, and Delete.
- Synchronization Object:** A section for selecting objects to synchronize, featuring a search bar and a list of databases such as 'baichao_increment', 'db001', 'db002', and various 'db_marisdb' instances.

Table 3-355 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-333 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> |

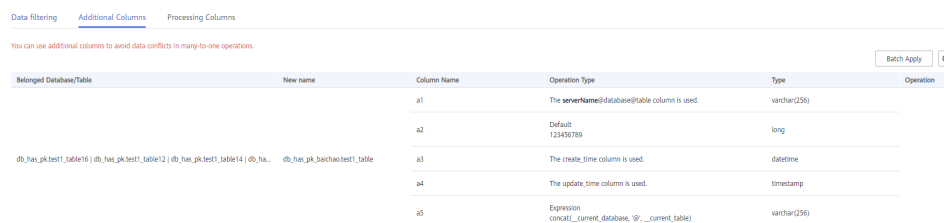
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: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

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

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

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

Figure 3-334 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**.

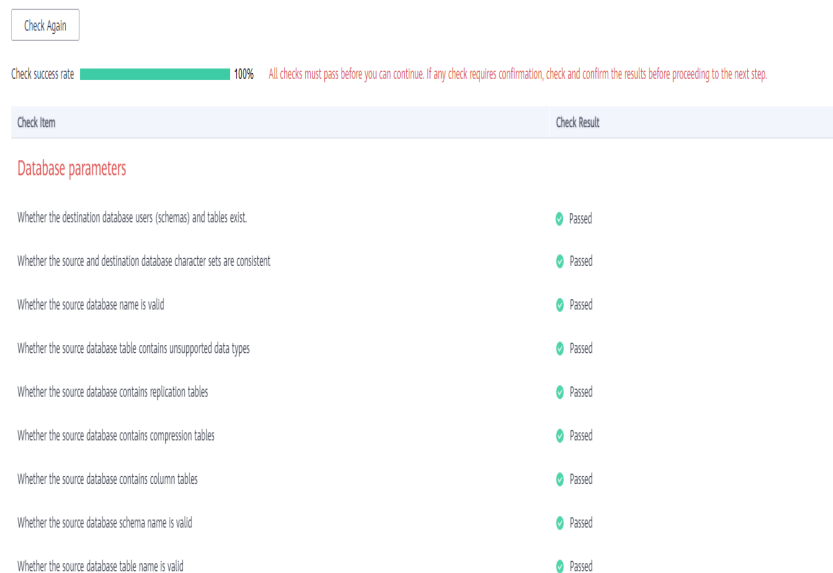
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Figure 3-335 Pre-check



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

Figure 3-336 Task startup settings

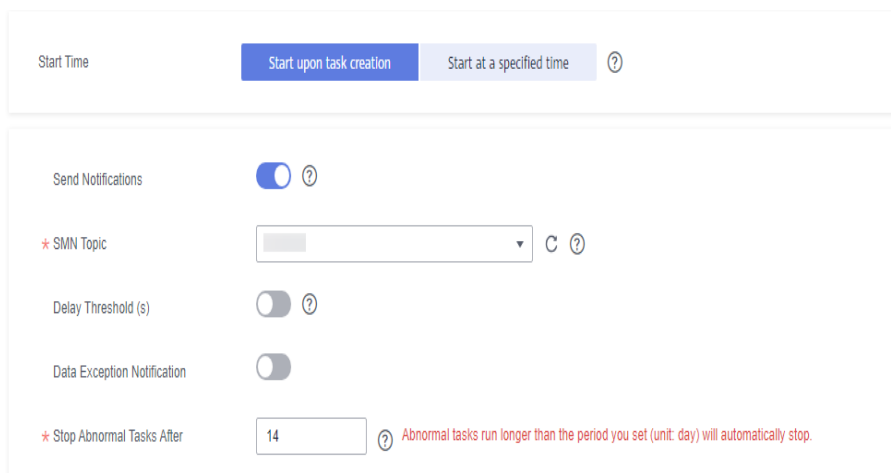



Table 3-356 Task startup settings

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

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

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

----End

3.33 From MariaDB to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-357 Supported databases

| Source DB | Destination DB |
|--|------------------------------------|
| <ul style="list-style-type: none"> • On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 • ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 • Other cloud MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 | GaussDB(for MySQL) Primary/Standby |

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

[Table 3-358](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 3-358 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> ● Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. ● Only MyISAM and InnoDB tables can be synchronized. ● Events and triggers cannot be synchronized. ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● 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. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur. |

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 3-359](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

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

Table 3-359 Database account permission

| Type | Incremental and Full+Incremental |
|---------------------------|---|
| Source database user | SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions) |
| Destination database user | The root account of GaussDB(for MySQL) has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-Level Comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

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 3-360 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - During the incremental synchronization, the binlog of the source MariaDB database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is a self-managed MariaDB 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 MariaDB instance, set the binlog retention period by following the instructions provided in RDS User Guide. - Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. - During an incremental synchronization, the source MariaDB database server_id must be set to a value ranging from 1 to 4294967296. ● Source database object requirements: <ul style="list-style-type: none"> - The source database names cannot contain non-ASCII characters, or the following characters: '<>/\'' - 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. ● Destination database parameter requirements: <ul style="list-style-type: none"> - Data cannot be synchronized from a newer version database to an older version database. - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> - 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. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database must be a primary/standby GaussDB(for MySQL) instance. - The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal. - The destination DB instance must have sufficient storage space. - If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. - The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys - If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. - If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. - If the sources and destinations are RDS instances, database mapping is required. - The source and destination databases cannot contain tables that have the same names but do not have primary keys. - 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"> - If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. - Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. - The source database does not support point-in-time recovery (PITR). - The destination database cannot be restored to a point in time when a full synchronization was being performed. - Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. - Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. - The partitioned table does not support column mapping. - After a task is created, the destination database cannot be set to read-only. - During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies. |

| Type | Restrictions |
|-----------------------------|---|
| | <ul style="list-style-type: none"> - If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |
| Full synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • Data inconsistency may occur when the MyISAM table is modified during a full synchronization. |
| 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. • During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • You can add additional objects during an incremental synchronization. • During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? |

| Type | Restrictions |
|----------------------------|---|
| Synchronization comparison | <ul style="list-style-type: none"> You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. |

Procedure

This section describes how to use DRS to configure a real-time synchronization task from MariaDB to GaussDB(for MySQL) over a public network.

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

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

- Task information description

Figure 3-337 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 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." The "Task Name" and "Description" fields have help icons. The "Description" field has a character count "0/256" at the bottom right.

Table 3-361 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

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

- Synchronization instance details

Figure 3-338 Synchronization instance details

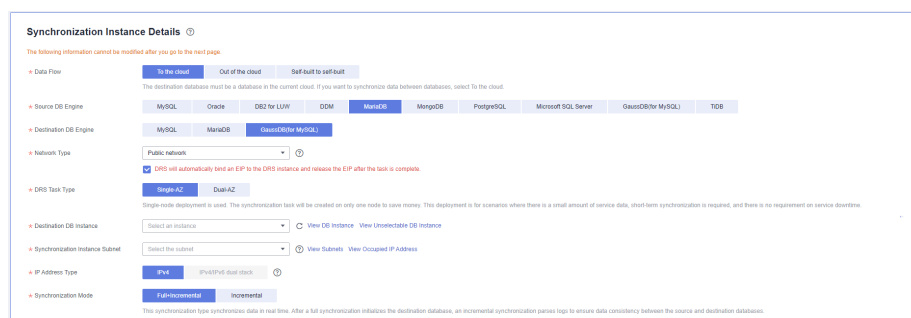


Table 3-362 Synchronization instance settings

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

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

| Parameter | Description |
|----------------------|--|
| Synchronization mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 3-339 Specifications




Table 3-363 Task type information

| Parameter | Description |
|-----------|--|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-340 AZ</p> |

- Enterprise Project and Tags

Figure 3-341 Enterprise projects and tags




Table 3-364 Enterprise Project and Tags

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

NOTE

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

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

- Source database information

Figure 3-342 Source database information

Source Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 3-365 Source database information

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 - 65535 |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting , Full synchronization , Incremental synchronization , or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk. |

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 3-343 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 3-366 Destination database settings

| Parameter | Description |
|-------------------|---|
| DB Instance Name | The GaussDB(for MySQL) instance you selected when creating the task. The parameter cannot be changed. |
| Database Username | The username for accessing the destination database. |
| Database Password | Password for logging in to the destination database. |

 **NOTE**

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

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

Figure 3-344 Synchronization objects

Flow Control Yes No ⓘ

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

Filter DROP DATABASE Yes No

Synchronize Normal index Incremental DDLs ⓘ

Synchronize DML Insert Update Delete ⓘ

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

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

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

Synchronization Object Tables Databases Import object file
Newly added tables in the database will be automatically synchronized to the destination, and deleted tables are automatically no longer synchronized.
 If any data in the source database changes, click the refresh button below.
 Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All ⓘ

Search the expanded database using regular expressions. 🔍

Select All ⓘ

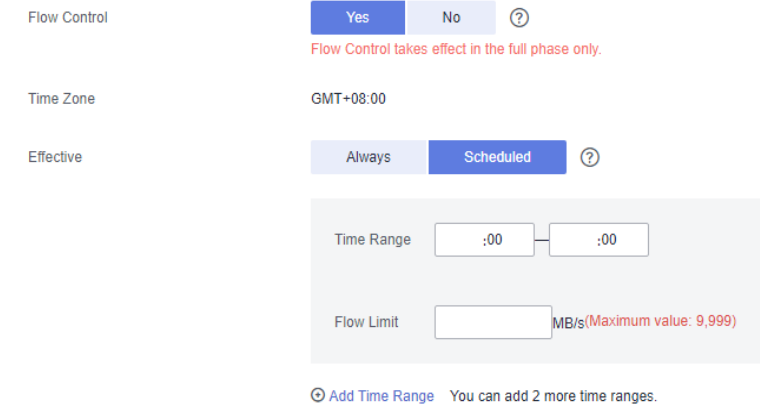
Search the expanded database using regular expressions. 🔍

db_0126 (New name: db_0126_new1) Edit database

➤ ➤


➤ ➤

Table 3-367 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-345 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> |

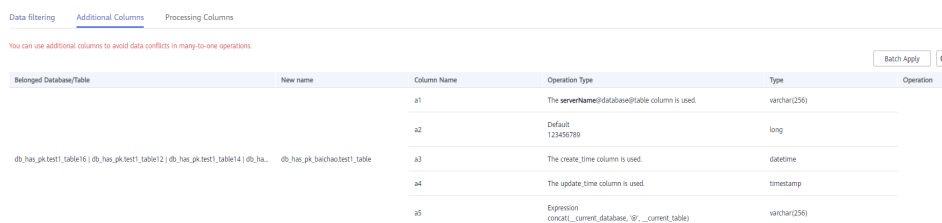
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: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

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

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

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

Figure 3-346 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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Figure 3-347 Pre-check



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

Figure 3-348 Task startup settings

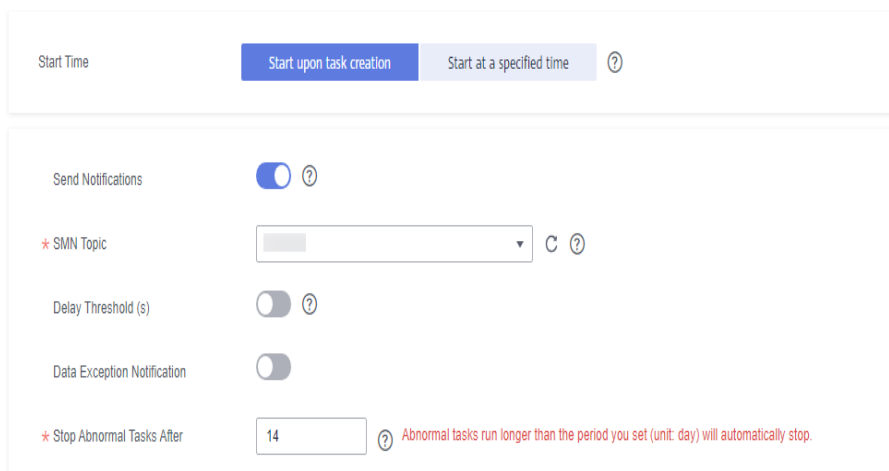



Table 3-368 Task startup settings

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

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

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

----End

3.34 From GaussDB(for MySQL) to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 3-369 Supported databases

| Source DB | Destination DB |
|------------------------------------|------------------------------------|
| GaussDB(for MySQL) Primary/Standby | GaussDB(for MySQL) Primary/Standby |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-370 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. |

| 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● The 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. ● 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 source database server_id must be set to a value ranging from 1 to 4294967296. ● 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 must be a primary/standby GaussDB(for MySQL) instance. ● 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 a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source and destination databases are the same GaussDB(for MySQL) instance, real-time synchronization without database mapping is not supported. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● 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 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. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● If you use additional columns and the number of columns in a single table exceeds 500, adding additional columns may lead to the number of columns in a table to reach an upper limit and cause the task to fail. ● Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. ● If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |

Procedure

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

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

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

- Task information description

Figure 3-349 Synchronization task information

Table 3-371 Task information

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

- Synchronization instance information

Figure 3-350 Synchronization instance information

Table 3-372 Synchronization instance settings

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

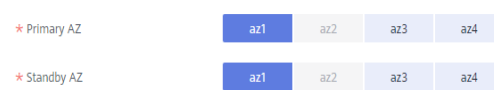
| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the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

- Task Type

Figure 3-351 Task type



Table 3-373 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 3-352 AZ</p>  |

- Enterprise Project and Tags

Figure 3-353 Enterprise projects and tags




Table 3-374 Enterprise Project and Tags

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

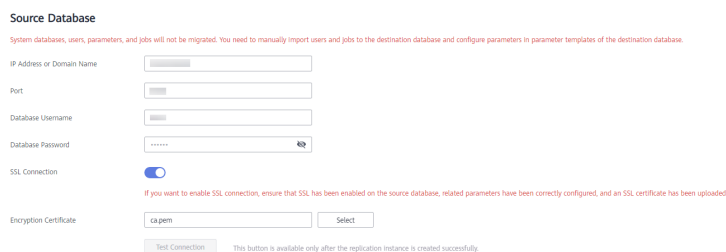
| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags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

 **NOTE**

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

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

Figure 3-354 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:

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

Encryption Certificate:

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

Table 3-375 Source database settings

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

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

 **NOTE**


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

Figure 3-355 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password 


 Test successful

Table 3-376 Destination database settings

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

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

Figure 3-356 Synchronization Mode

Flow Control: Yes No ⓘ

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

Read on Standby Node: ⓘ

Filter DROP DATABASE: Yes No

Synchronize: Normal index ⓘ

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

Synchronize DDLs: Default Custom ⓘ

During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN.

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

Synchronization Object: Tables Databases Import object file

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

Select All

Search the expanded database using regular expressions.

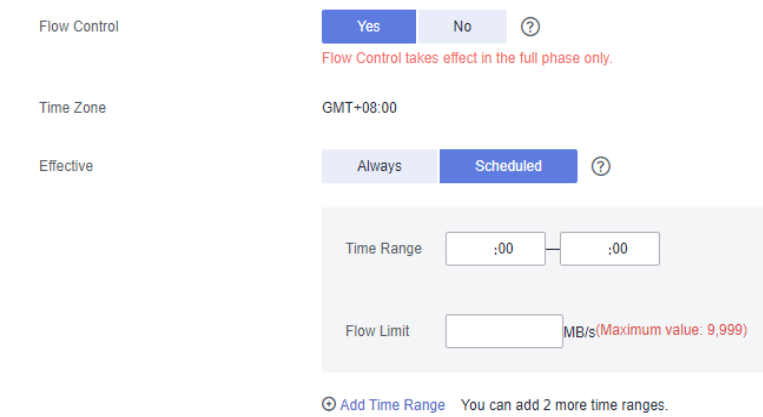
- db_atonchange_002_three database
- db_atonchange_002_two database

Select All

Search the expanded database using regular expressions.

- db_atonchange_002 Edit database


Table 3-377 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 3-357 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| Read on Standby Node | <p>If the source in Step 3 is an optional instance, data can be read from the standby database. You can enable or disable the function.</p> <ul style="list-style-type: none"> ● If enabled, full data is read from the standby database, and incremental binlogs are obtained from the primary database. <ul style="list-style-type: none"> – After this function is enabled, the primary and standby databases at the source end must be available, and the synchronization between the primary and standby databases is normal. If there are tables without primary keys in the synchronization object, the source database user must have the RELOAD permission. – After DDLs are executed on the primary database at the source end, the standby database disconnects the synchronization between the primary and standby databases. As a result, full synchronization of the DRS task may fail. If the table structure is not modified, you can resume the task. – If the primary/standby role at the source end changes, full synchronization of the DRS task may fail. In this case, you can resume the synchronization to restore the task. – Data can be read from the standby database only when the source is an optional instance. If the source is a self-managed database, the IP address cannot be entered. ● If disabled, both full data and incremental binlogs are obtained from the primary database. |

| Parameter | Description |
|-------------------------------|--|
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Whether to synchronize normal indexes.</p> <p>By default, DRS synchronizes the primary key or unique index. A normal index refers to an index other than the primary key or unique index. If you select Normal index, all indexes will be synchronized. If you do not select normal index, only the primary key or unique index will be synchronized.</p> |
| Start Point | <p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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>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. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

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

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

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

Figure 3-358 Processing data

| Belonged Database/Table | New name | Column Name | Operation Type | Type | Operation |
|-------------------------|--------------|-------------|----------------|------|-----------|
| db1_jack_011 | db1_jack_011 | -- | -- | -- | Add |
| db1_1ae2 | db1_1ae2 | -- | -- | -- | Add |
| db1_1ae3 | db1_1ae3 | -- | -- | -- | Add |

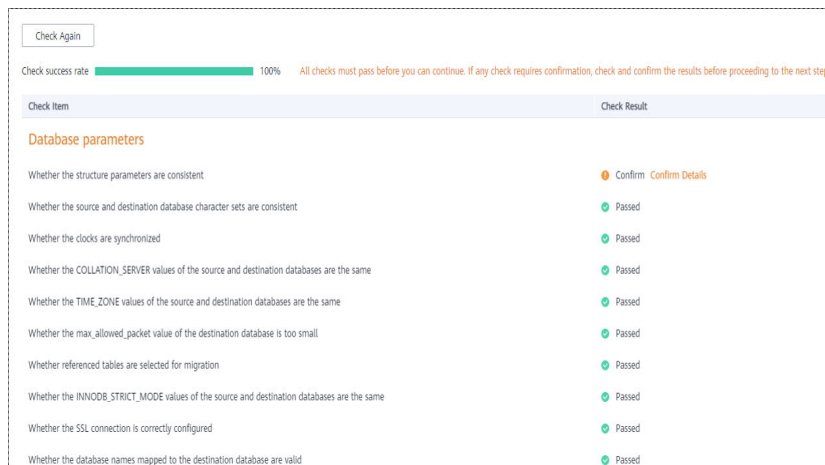
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

Figure 3-359 Pre-check



NOTE

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

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

Figure 3-360 Task startup settings

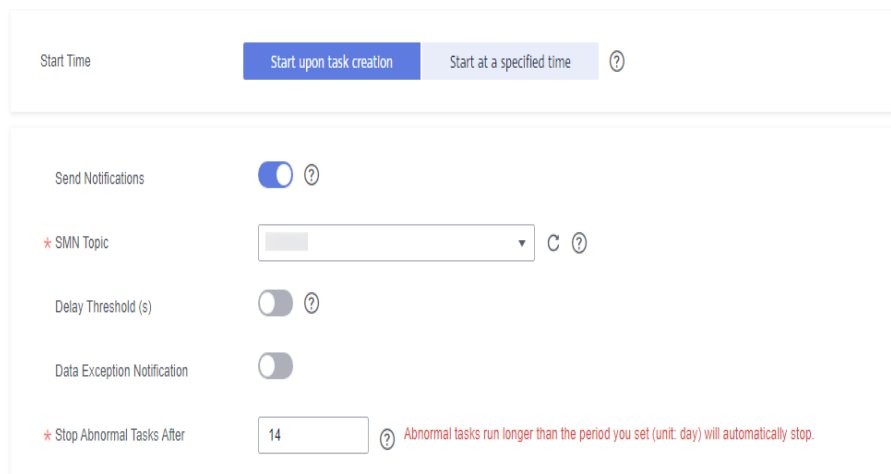



Table 3-378 Task startup settings

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

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

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

----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 <p>NOTE The destination database version must be the same as or later than the source database version.</p> |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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. Otherwise, all synchronization tasks fail.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) 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"><li data-bbox="587 344 1369 472">• The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT<li data-bbox="587 488 1414 723">• The destination database user must have the following permissions: SELECT, CREATE, INDEX, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● Only tables, primary key indexes, unique indexes, common indexes, 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. ● Braces ({}), in a database or table name in the source database must be used in pairs. Otherwise, the table structure synchronization may fail. After data mapping, if braces ({}), in a database or name mapped to the destination database are not in pairs, the table structure may fail to be synchronized. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|-----------------|---|
| Source database | <ul style="list-style-type: none"> ● The source database names cannot contain non-ASCII characters, or the following characters: '<>/' ● 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● The source database cannot be a read replica. |

| Type | Restrictions |
|----------------------|--|
| Destination database | <ul style="list-style-type: none"> • Data cannot be synchronized from a newer version database to an older version database. • The destination database must have sufficient disk space. • The character set of the destination database must be the same as that of the source database. • The time zone of the destination database must be the same as that of the source database. • If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. • During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. • If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails. • The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see Syntax Differences Between MySQL Versions. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● If the sources and destinations are RDS instances, database mapping is required. ● The source and destination databases cannot contain tables that have the same names but do not have primary keys. ● If the source 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. ● If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Restrictions |
|------|--|
| | <p>requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • The source database does not support point-in-time recovery (PITR). • The destination database cannot be restored to a point in time when a full synchronization was being performed. • Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. • Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. • The 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. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • DDL operations are not supported during full synchronization. • During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • You can add additional objects during an incremental synchronization. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |

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

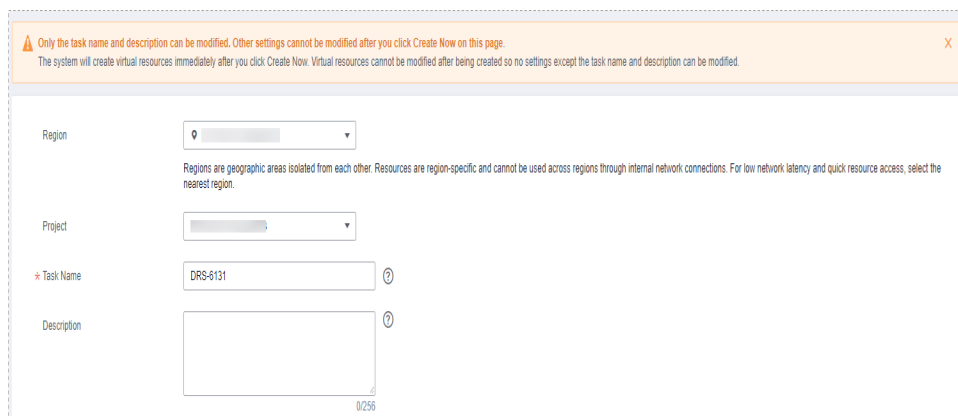


Table 4-3 Task information

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

- Synchronization instance details

Figure 4-2 Synchronization instance details

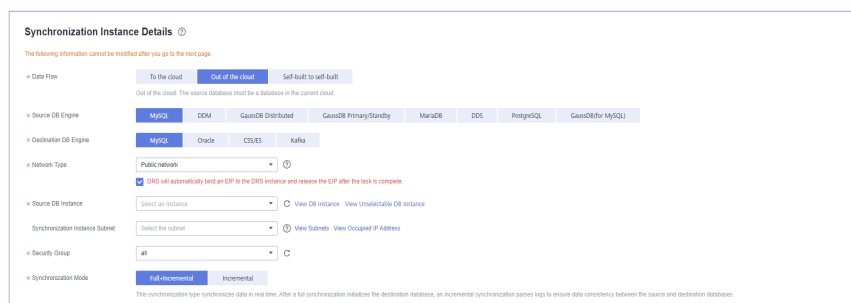


Table 4-4 Synchronization instance settings

| Parameter | Description |
|-----------|--|
| Data Flow | Select Out of the cloud . The source database is a database on the current cloud. |

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

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <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. |

- Task type

Figure 4-3 Task type



Table 4-5 Task type information

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

- Enterprise Project and Tags

Figure 4-4 Enterprise projects and tags

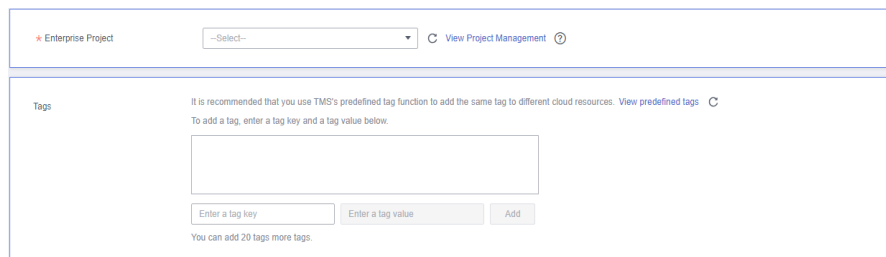


Table 4-6 Enterprise Project and Tags

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

NOTE

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

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

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


- Source database information

Figure 4-5 Source database information

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 disabled, your data may be at risk. |

 **NOTE**

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

- Destination database information

Figure 4-6 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 4-8 Destination database settings

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

 **NOTE**

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

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

Figure 4-7 Synchronization mode

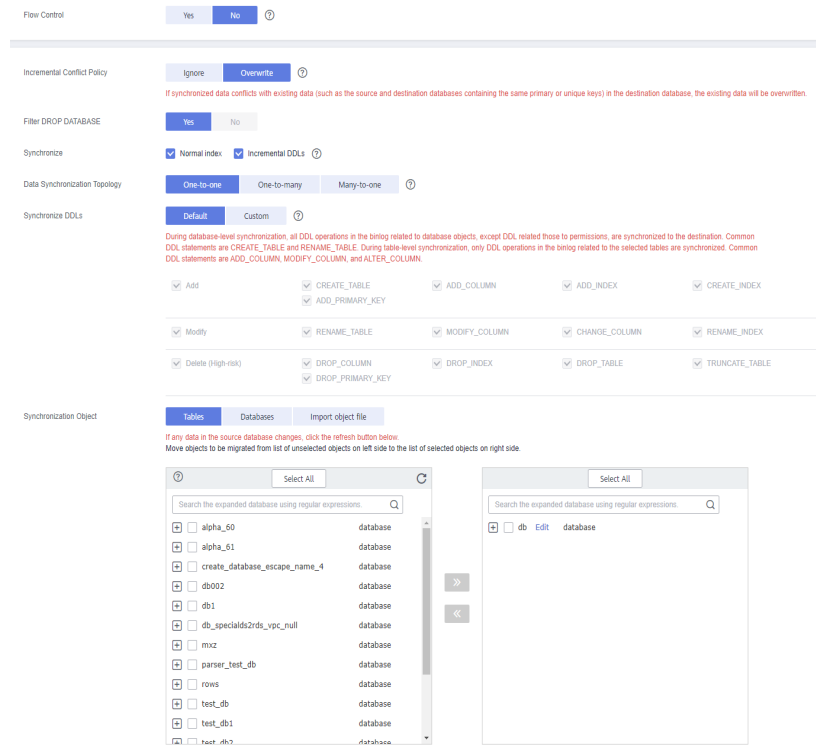
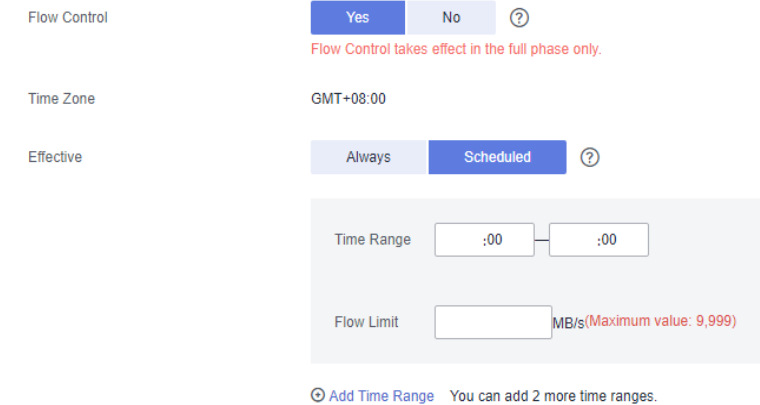



Table 4-9 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize 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 position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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> |

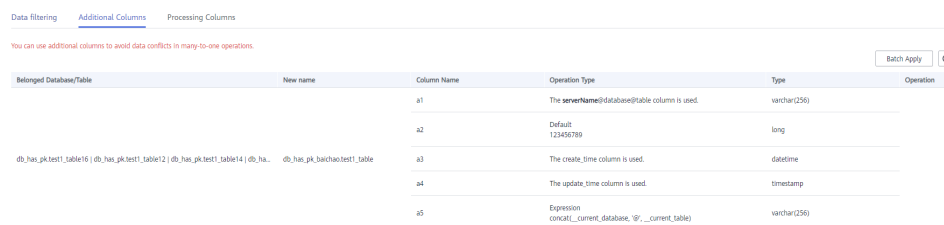
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: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

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

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

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

Figure 4-9 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, '@', current_table) | varchar(256) | |

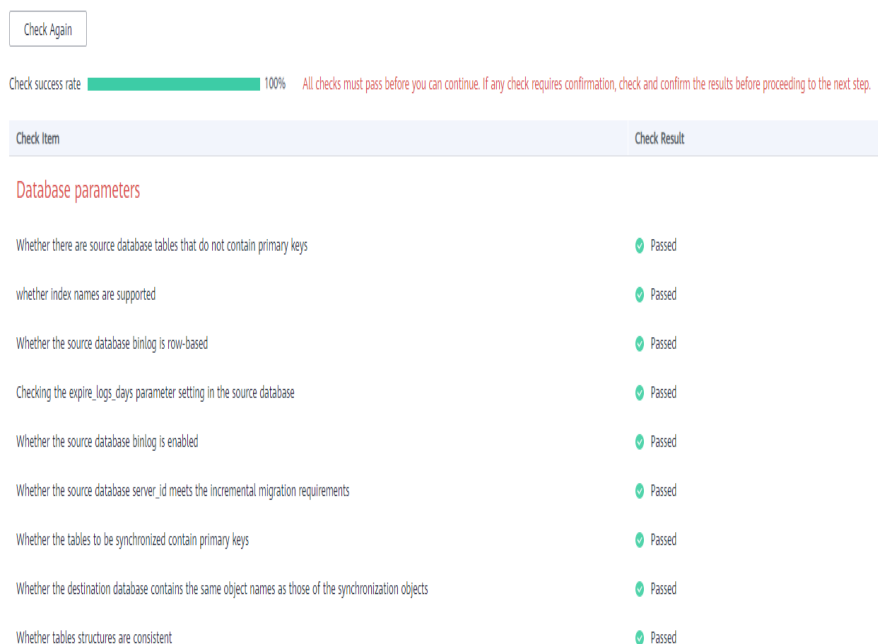
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Solutions to Failed Check Items](#) 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-11 Task startup settings

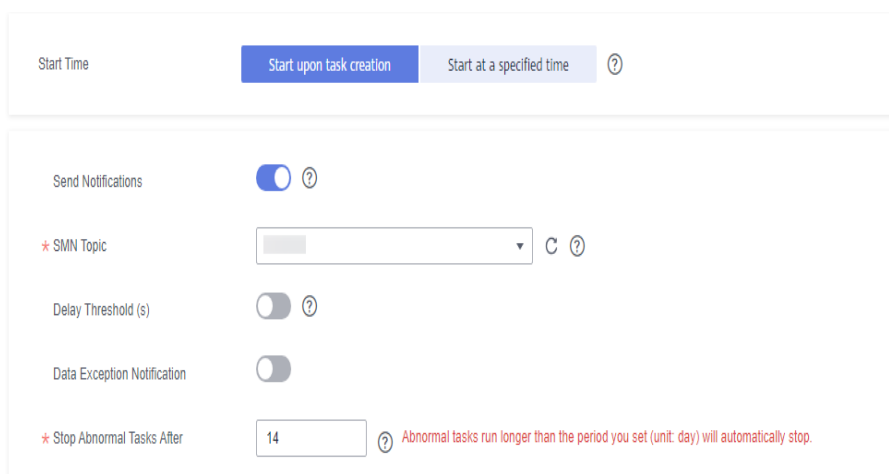



Table 4-10 Task startup settings

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

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

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

----End

4.2 From MySQL to Kafka

Supported Source and Destination Databases

Table 4-11 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none"> • RDS for MySQL | <ul style="list-style-type: none"> • Kafka |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - Tables to be synchronized without a primary key may be locked for 3s.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

- If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-12 Precautions

| Type | Restrictions |
|------------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, SHOW VIEW, EVENT, REPLICATION SLAVE, and REPLICATION CLIENT. |
| Synchronization object | <ul style="list-style-type: none"> • During full synchronization, tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized, but events and triggers cannot be synchronized. During incremental synchronization, only table data and DDLs can be synchronized. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|---|
| Source database | <ul style="list-style-type: none"> ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● The database and table names in the source database cannot contain non-ASCII characters, or 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. ● You are advised to set auto.create.topics.enable of Kafka to false. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the data types are incompatible, the synchronization may fail. ● If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● A real-time synchronization task may fail due to the change of the username and password of the source database. You need to rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization. ● If the source database port is changed during data synchronization, the synchronization task fails. If the destination database port is wrong, DRS automatically changes the port to the correct one, and then you need to retry the synchronization task. Generally, do not modify the port number during synchronization. ● If a real-time synchronization task fails as the IP address is changed, the system automatically changes the IP address to the correct one. Then, you need to retry the task to continue the synchronization. Therefore, changing the IP address is not recommended. ● If a full synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the identifier field in the Kafka data for data deduplication. (The shard ID must be unique.) ● During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database does not support point-in-time recovery (PITR). ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. |

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

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 with the value "DRS-6131"), and "Description" (a text area). 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." The bottom right corner of the form area shows "0/256".

Table 4-13 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

| 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 4-13 Synchronization instance details

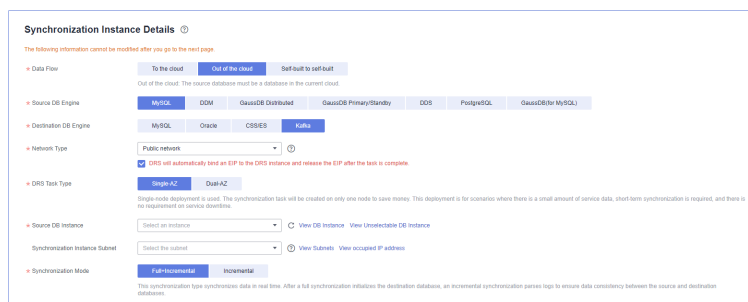


Table 4-14 Synchronization instance settings

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

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

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

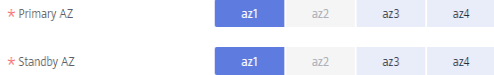
- Task Type

Figure 4-14 Task type



Table 4-15 Task type information

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

| Parameter | Description |
|-----------|--|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-15 AZ</p>  |

- Enterprise Project and Tags

Figure 4-16 Enterprise projects and tags




Table 4-16 Enterprise Project and Tags

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

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

 **NOTE**

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


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

Figure 4-17 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ 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. |


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

 **NOTE**

The username and password of the source database are encrypted and stored in 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**.

The parameters on the **Set Synchronization Task** page vary according to the synchronization mode selected in **Step 2**. The **Full+Incremental** synchronization mode is used as an example in **Figure 4-19**.

Figure 4-19 Synchronization mode

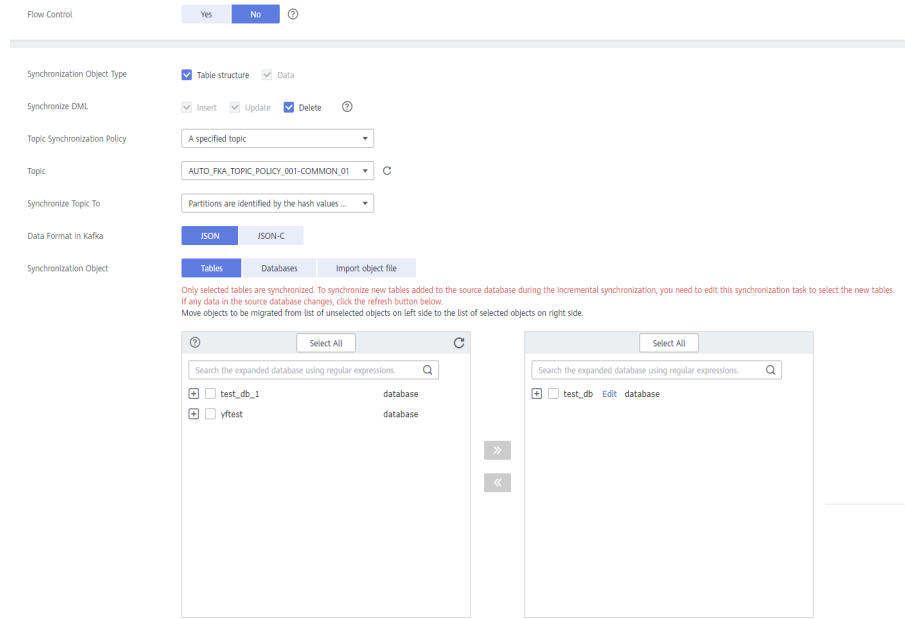
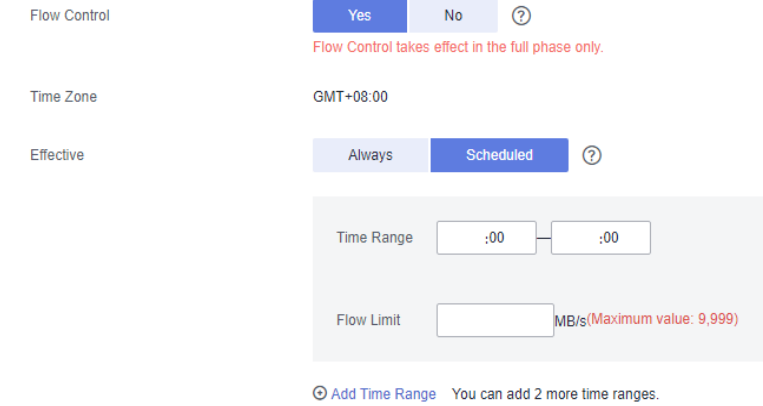



Table 4-19 Synchronization Object

| Parameter | Description |
|-----------------------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. |
| Synchronization Object Type | <p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> |
| Synchronize DML | <p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> |

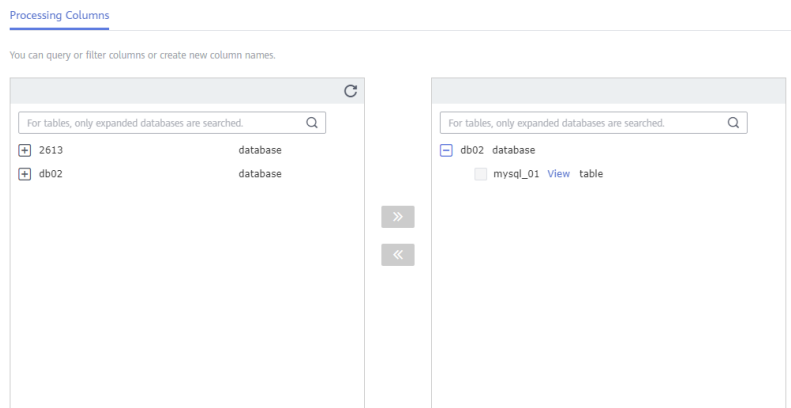
| Parameter | Description |
|------------------------------|---|
| Start Point | <p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| Topic Synchronization Policy | <p>Topic synchronization policy. You can select A specific topic or Auto-generated topics.</p> |
| Topic | <p>Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic.</p> |
| Topic Name Format | <p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>Only variables database and tablename are supported. The other characters must be constants. Replace \$database\$ with the database name and \$tablename\$ with the table name.</p> <p>For example, if this parameter is set to \$database\$-\$tablename\$ and the database name is db1, and the table name is tab1, the topic name is db1-tab1. If DDL statements are synchronized, \$tablename\$ is empty and the topic name is db1.</p> |
| Number of Partitions | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p> |
| Replication Factor | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p> |

| Parameter | Description |
|------------------------|--|
| Synchronize Topic To | <p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. • 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 from MySQL to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. This option is available only when Synchronization Mode is set to Incremental in Step 2. Only whitelisted users can use the Avro option. To use this option, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket. • JSON: JSON message format, which is easy to interpret but takes up more space. • JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

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

Figure 4-21 Processing data



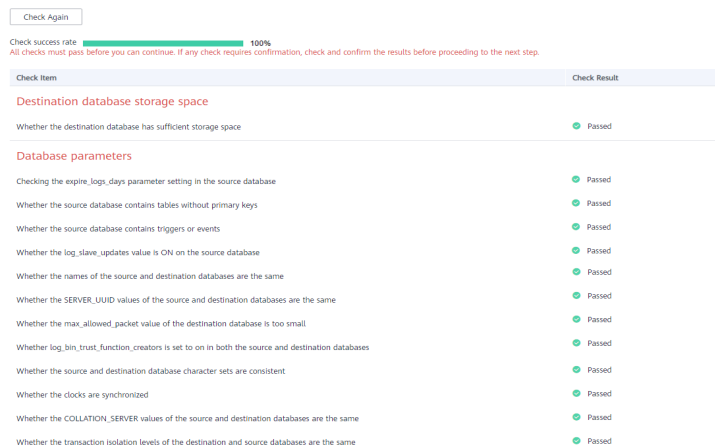
Step 6 On the **Check Task** page, check the synchronization task.

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

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

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

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-23 Task startup settings





Table 4-20 Task startup settings

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

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

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

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

----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 |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-22 Precautions

| Type | Restrictions |
|------------------------|--|
| Database permissions | <p>Minimum permission requirements for full plus incremental synchronization:</p> <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT • The destination database user must have the following permissions: READ and WRITE |
| Synchronization object | <ul style="list-style-type: none"> • The table data can be synchronized. • Databases, views, indexes, constraints, functions, stored procedures, triggers, and events cannot be synchronized. • The system database and event status cannot be synchronized. • Tables whose primary keys are of the FLOAT type cannot be synchronized. • Tables that do not have primary keys cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|---|
| Source database | <ul style="list-style-type: none"> ● The source database 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails. ● The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail. ● The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail. ● The source database cannot be a read replica. |
| Destination database | <ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored to a point in time when a full synchronization was being performed. ● If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● 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. |

| 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. • 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. 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 4-23 Task information

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

- Synchronization instance details

Figure 4-25 Synchronization instance details

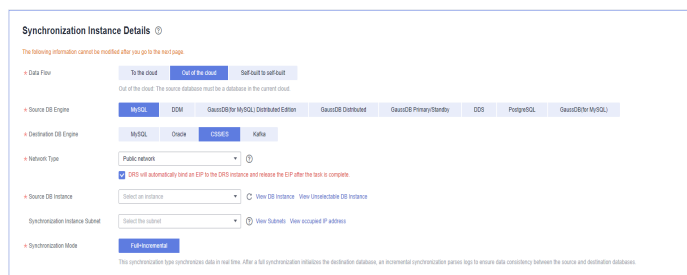


Table 4-24 Synchronization instance settings

| Parameter | Description |
|-----------------------|----------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select MySQL . |
| Destination DB Engine | Select CSS/ES . |

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

- Task type

Figure 4-26 Task type




Table 4-25 Task type information

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

- Enterprise Project and Tags

Figure 4-27 Enterprise projects and tags




Table 4-26 Enterprise Project and Tags

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

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

 **NOTE**

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


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

Figure 4-28 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. |

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

 **NOTE**

The username and password of the source database are encrypted and stored in 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. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

| Parameter | Description |
|----------------|---|
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. Only .cer and .pem certificates are supported. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-30 Synchronization mode

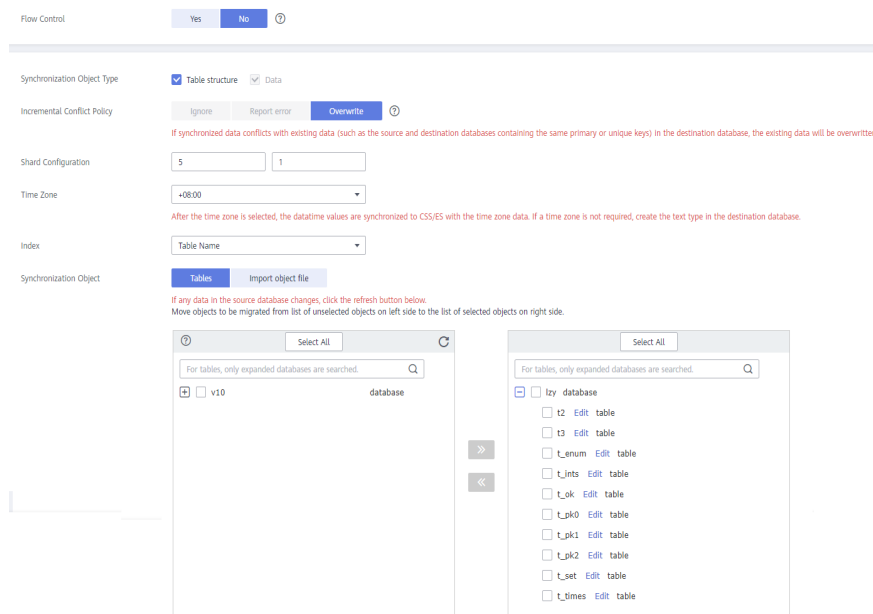
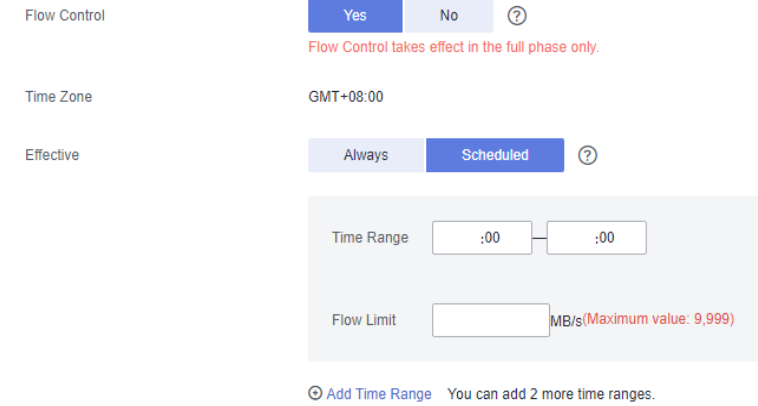



Table 4-29 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy only applies to incremental synchronization. The default value is Overwrite. The conflict in the full synchronization phase is ignored by default.</p> |
| Shard Configuration | <p>Configure the number of primary shards and the number of replicas. The default number of primary shards is 5, and the default number of shard replicas is 1.</p> |
| Time Zone | <p>After a time zone is selected, the datetime values are synchronized to CSS/ES with the time zone data.</p> |
| Index Name | <ul style="list-style-type: none"> • Table Name The index name created in the target Elasticsearch instance is the same as the table name. • Database Name+Table Name. The name of the index created in the target Elasticsearch instance is <i>DatabaseName_TableName</i>. |

| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> You can select objects by importing a file. For details, see Importing Synchronization Objects. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. If you select Tables, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the limit is exceeded, you can add synchronization objects in batches by editing the synchronization objects. For details, see Editing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering** or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 4-32 Processing data

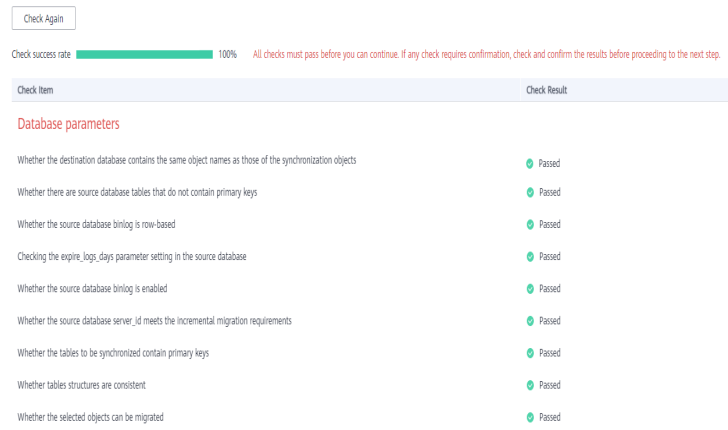


- Step 6** On the **Check Task** page, check the synchronization task.
- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-34 Task startup settings

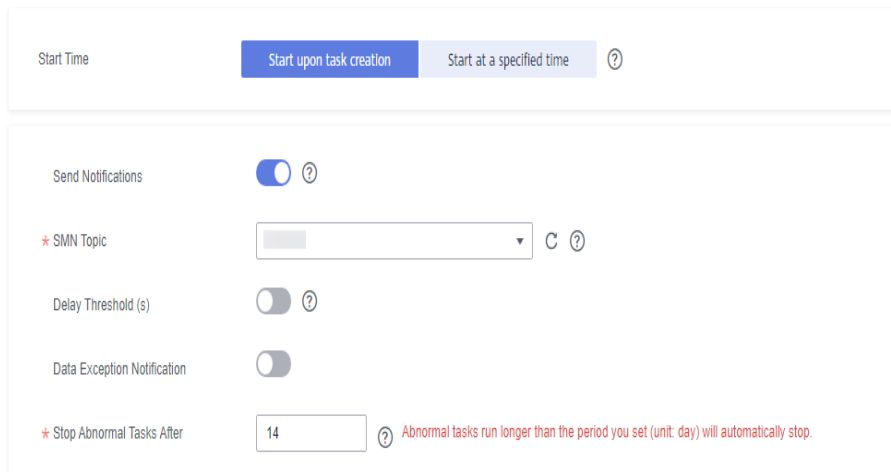



Table 4-30 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----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 |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- Data-level comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) 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, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. • The destination database user must have the following permissions: ALTER ANY INDEX, ALTER ANY TABLE, ALTER SESSION, ANALYZE ANY, COMMENT ANY TABLE, CREATE ANY INDEX, CREATE ANY TABLE, CREATE SESSION, DELETE ANY TABLE, DROP ANY TABLE, INSERT ANY TABLE, SELECT ANY TABLE, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, UPDATE ANY TABLE, and RESOURCE roles. |

| Type | Restrictions |
|------------------------|--|
| Synchronization object | <ul style="list-style-type: none"> ● Full synchronization supports the synchronization of data, table structures, and indexes. ● Incremental synchronization supports only data synchronization. ● Geography data types such as geometry, geometrycollection, linestring, multilinestring, multipoint, point and polygon are not supported. ● Views, constraints, functions, stored procedures, triggers, and events cannot be synchronized. ● The system database and event status cannot be synchronized. ● The destination Oracle database does not support empty strings, so the object to be synchronized cannot contain empty strings. ● The maximum number of columns supported by the source MySQL database is 1017, while the maximum number of columns supported by the destination Oracle database is 1000. Therefore, the number of columns of objects to be synchronized cannot exceed 1000. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|-----------------|--|
| Source database | <ul style="list-style-type: none"> • 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. • During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. • Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. • GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. • The source database does not support the mysql binlog dump command. • The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails. • The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail. • The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail. • If the source MySQL database version is 8.0, do not set lower_case_table_names to 0. • The source database cannot be a read replica. |

| 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. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● The NOT NULL constraint of MySQL supports empty strings, while the NOT NULL constraint of Oracle does not. During a synchronization, if an empty string exists in a NOT NULL constraint field, delete the NOT NULL constraint from the destination Oracle database. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● If the default value of the time field in the source database is all 0s, it will be converted to 1970-01-01 00:00:00. ● If the precision of source database decimal data type exceeds 38, the data will be truncated because the maximum precision of the destination database number data type is 38. ● The varbinary, binary, and timestamp columns in the source database cannot contain primary key or unique constraints. ● If the length of the varchar data in the source database is greater than or equal to 667 characters, the varchar type will be converted to the clob type in the Oracle database. ● 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. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • 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. • 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 information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-36 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 | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The RDS for MySQL instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full +Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. |

- Task type

Figure 4-37 Task type



Table 4-35 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-38 Enterprise projects and tags




Table 4-36 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.


Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-39 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. |

| Parameter | Description |
|----------------|--|
| SSL Connection | <p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 4-40 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 | <p>The IP address or domain name of the destination database.</p> <p>NOTE</p> <p>For a RAC cluster, use a scan IP address to improve access performance.</p> |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |

| Parameter | Description |
|-------------------|--|
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **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

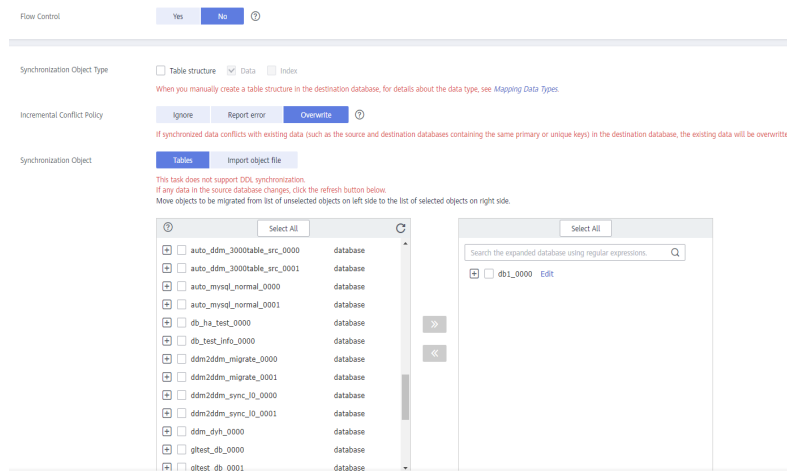
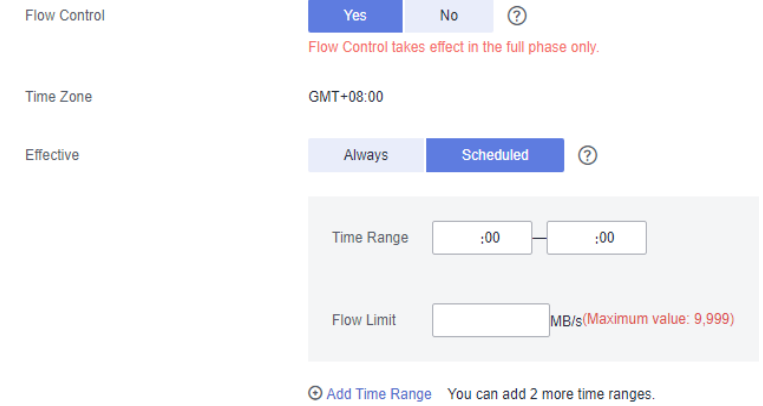



Table 4-39 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization 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 Always. 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>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <p>For details about how to import an object file, see Importing Synchronization Objects.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

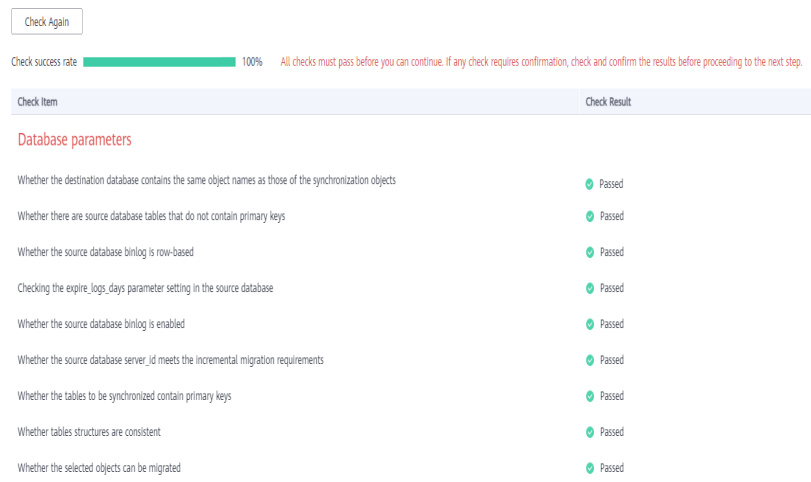
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-44 Task startup settings





Table 4-40 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.5 From MySQL to MariaDB

Supported Source and Destination Databases

Table 4-41 Supported databases

| Source DB | Destination DB |
|---|--|
| <ul style="list-style-type: none"> • RDS for MySQL 5.6, 5.7, and 8.0 | <ul style="list-style-type: none"> • On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 • ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 • Other cloud MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 <p>NOTE</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6, you are advised to select MariaDB 10.0, 10.1, or later as the destination database. • If the source database version is MySQL 5.7, you are advised to select MariaDB 10.2, 10.3, 10.4, or later as the destination database. • If the source database version is MySQL 8.0, you are advised to select MariaDB 10.5 as the destination database. |

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-42](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 4-42 Database account permission

| Type | Full Synchronization | Full+Incremental Synchronization and Incremental Synchronization |
|---------------------------|---|---|
| Source database user | The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT | The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT |
| Destination database user | The user must have the following minimum permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES | |

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Supported Synchronization Objects

Table 4-43 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-43 Supported synchronization objects

| Type | Precautions |
|-------------------------|--|
| Synchronization objects | <ul style="list-style-type: none"> ● Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. ● Only MyISAM and InnoDB tables can be synchronized. ● Events and triggers cannot be synchronized. ● The table structure of the synchronization object in the source database cannot contain the table identifier ENCRYPTION. ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● 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. ● 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling suggestions, see [Pre-check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 4-44](#).

Table 4-44 Precautions

| Type | Restrictions |
|-----------------|--|
| 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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">• 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 and MariaDB system databases) 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. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the source and destination database versions are different, syntax compatibility issues may occur due to feature differences between the source and destination database versions. For details, see What Are Syntax Differences Between MySQL or MariaDB Versions? ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● If the source database contains any table without a primary key defined, the table-level many-to-one synchronization is not supported. ● 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 MySQL database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. |

| 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. ● 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. ● DDL operations are not supported during full synchronization. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. ● During table-level synchronization, online DDL can be used for incremental synchronization. For details, see Does DRS Support Online DDL for Real-Time Synchronization? |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • You can add additional objects during an incremental synchronization. • During table-level synchronization, in the many-to-one scenario where an additional column is set as the source column on the data processing page, if there is a mapped table in the destination database, delete the table or clear data in the table in the destination database. Otherwise, the composite primary key will not be created by adding additional columns. This will cause data conflicts during data synchronization. If the data conflicts are ignored, there may be data inconsistencies. • If the source database version is MySQL 8.0, some collation character sets (such as utf8mb4_0900_as_ci, utf8mb4_0900_as_cs, utf8mb4_0900_bin and utf8mb4_cs_0900_ai_ci that support Unicode 9.0) are not supported by the destination database. • Due to engine and version differences, the following MySQL functions are not supported in MariaDB: MBRCOVEREDBY, ST_BUFFER_STRATEGY, ST_GeoHash, ST_IsValid, ST_LatFromGeoHash, ST_LongFromGeoHash, ST_PointFromGeoHash, ST_SIMPLIFY, ST_VALIDATE, (8.0)JSON_ARRAYAGG, JSON_OBJECTAGG, JSON_PRETTY, JSON_STORAGE_FREE, JSON_STORAGE_SIZE and JSON_TABLE. • The destination MariaDB database does not support VALIDATION during partition exchange. Do not use the verification syntax when executing partition exchange SQL statements in the source MySQL database. For example: ALTER TABLE t1 EXCHANGE PARTITION p0 WITH TABLE t2 WITH VALIDATION; • Do not write characters such as \n, \t, and \r to the JSON data of the source MySQL database during incremental synchronization. Otherwise, there may be data inconsistencies. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |

Procedure

This section describes how to synchronize data out of the cloud from an RDS for MySQL database to a MariaDB database. To configure other storage engines, you can refer to the following procedures.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-45 Synchronization task information




Table 4-45 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 4-46 Synchronization instance information

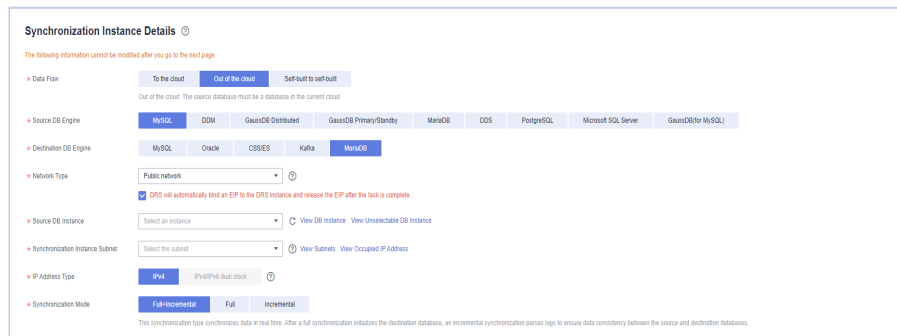


Table 4-46 Synchronization instance settings

| Parameter | Description |
|-----------------------|--|
| Data Flow | Select Out of the cloud . The source database is a database on the current cloud. |
| Source DB Engine | Select MySQL . |
| Destination DB Engine | Select MariaDB . |
| Network Type | Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for data synchronization between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| DRS Task Type | Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none">- Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.- Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task . |
| Source DB Instance | The RDS for MySQL instance you created. |

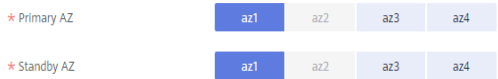
| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the source DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- AZ

Figure 4-47 AZ



Table 4-47 Task AZ

| Parameter | Description |
|-----------|--|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-48 AZ</p>  |

- Enterprise Project and Tags

Figure 4-49 Enterprise projects and tags

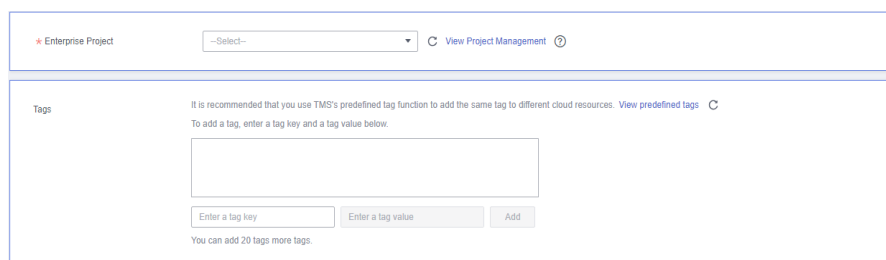


Table 4-48 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.


- Source database information

Figure 4-50 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 4-49 Source database information

| Parameter | Description |
|------------------|---|
| DB Instance Name | The RDS for MySQL instance selected during synchronization task creation. This parameter cannot be changed. |

| Parameter | Description |
|-------------------|--|
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>If SSL connection is required, enable SSL on the source database, ensure that related parameters have been correctly configured, and upload an SSL certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk. |


 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

- Destination database information

Figure 4-51 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 4-50 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Database Username | The username for accessing the destination database. |

| Parameter | Description |
|-------------------|---|
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - Only .cer and .pem certificates are supported. - If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 4-52 Synchronization Mode

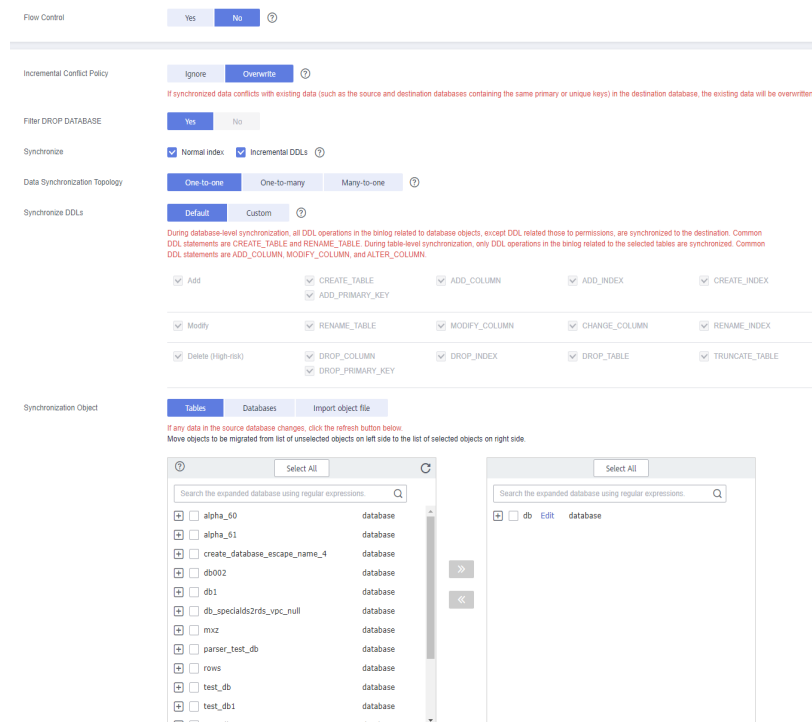
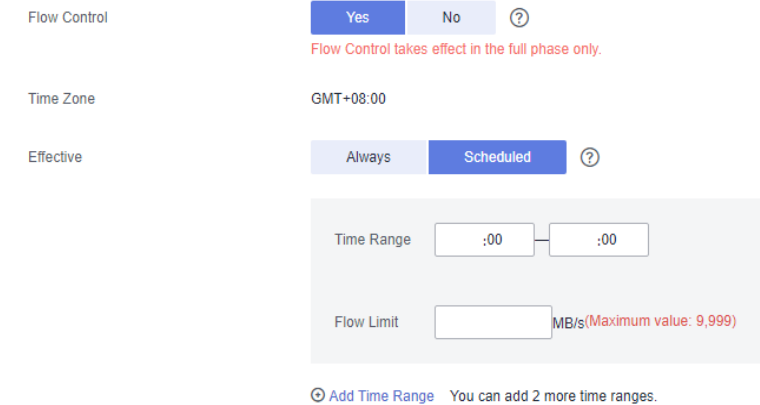



Table 4-51 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-53 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-------------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize 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 position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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> |

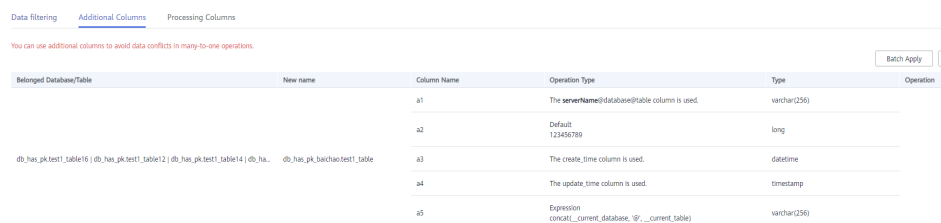
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: 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. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables. |

| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Mapping Object Names. <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (<code>_</code>), or the mapped database table name contains hyphens (<code>-</code>) and number signs (<code>#</code>), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 4-54 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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-55 Task startup settings

Table 4-52 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.6 From DDM to MySQL

Supported Source and Destination Databases

Table 4-53 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none">• DDM instances | <ul style="list-style-type: none">• On-premises MySQL databases• MySQL databases on an ECS• MySQL databases on other clouds• RDS for MySQL |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-54 Precautions

| Type | Restrictions |
|------------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The destination database user must have the following permissions: 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. – Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. ● 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. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> ● During an 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. - Online DDL is not supported. - Tables can be created. For example: <pre>create table `ddl_test` (id int, c1 varchar(25), primary key(id)); create table `ddl_test_gho` like `ddl_test`;</pre> - Tables can be renamed. Both the source and destination tables must be selected. For example: <pre>rename table `ddl_test` to `ddl_test_new`;</pre> - Columns in a table can be added and modified, but cannot be deleted. For example: <pre>alter table `ddl_test` add column `c2` varchar(25); alter table `ddl_test` modify column `c1` varchar(50); alter table `ddl_test` alter c1 set default '***';</pre> - Table indexes can be modified. For example: <pre>alter table `ddl_test` drop primary key; alter table `ddl_test` add primary key(id); alter table `ddl_test` add index `ddl_test_uk`(id); alter table `ddl_test` drop index `ddl_test_uk`;</pre> - In table-level synchronization, you can add columns, modify columns, and add primary keys and normal indexes. - 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. ● When editing a task, do not import a lot of data to the newly-added table. You are advised to edit the task during off-peak hours. |

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-56 Synchronization task information

The screenshot shows a form titled 'Create Synchronization Instance'. 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 this, there are four main sections:

- Region:** A dropdown menu with a search icon and a list of regions. Below it, a note 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.'
- Project:** A dropdown menu.
- Task Name:** A text input field containing 'DRS-B131' and a help icon. A note indicates it must start with a letter and consist of 4 to 50 characters, allowing only letters, digits, hyphens, and underscores.
- Description:** A text area containing a help icon and a character count '0/256'.

Table 4-55 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 4-57 Synchronization instance details

The screenshot shows the 'Synchronization Instance Details' configuration page. It includes several sections:

- Data Flow:** Three tabs: 'To the cloud', 'Out of the cloud' (selected), and 'Self-built to self-built'.
- Source DB Engine:** A row of buttons for MySQL, EDM (selected), GaussDB Distributed, GaussDB Primary/Standby, MariaDB, DDS, PostgreSQL, and GaussDB(for MySQL).
- Destination DB Engine:** A row of buttons for MySQL, Oracle, and Kafka.
- Network Type:** A dropdown menu set to 'Public network'.
- Source DB Instance:** A dropdown menu with a 'Select an instance' button and links for 'View DB Instance' and 'View Unselectable DB Instance'.
- Synchronization Instance Subnet:** A dropdown menu with a 'Select the subnet' button and a link for 'View Subnets'.
- Synchronization Mode:** Three tabs: 'Full-incremental' (selected), 'Full', and 'Incremental'. A note explains that 'Full-incremental' synchronizes data in real time and uses incremental synchronization for consistency.
- Source DB Instance Quantity:** A numeric input field set to '2'.

Table 4-56 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select DDM . |
| Destination DB Engine | Select MySQL . |
| Network Type | Public network is used as an example. Available options: VPC , Public network , and VPN or Direct Connect <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The DDM instance you created. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |

| Parameter | Description |
|-----------------------------|--|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |
| Source DB Instance Quantity | Specifies the number of DB instances bound to the source DDM database. |

- Task type

Figure 4-58 Task type



Table 4-57 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |

| Parameter | Description |
|-----------|---|
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-59 Enterprise projects and tags




Table 4-58 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-60 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 4-59 Source database settings

| Parameter | Description |
|-------------------|---|
| DB Instance Name | The DDM instance you selected when you create a synchronization task. The instance name cannot be changed. |
| Database Username | The username for accessing the source database. |
| Database Password | The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted. |

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 4-61 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-60 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-62 Synchronization mode

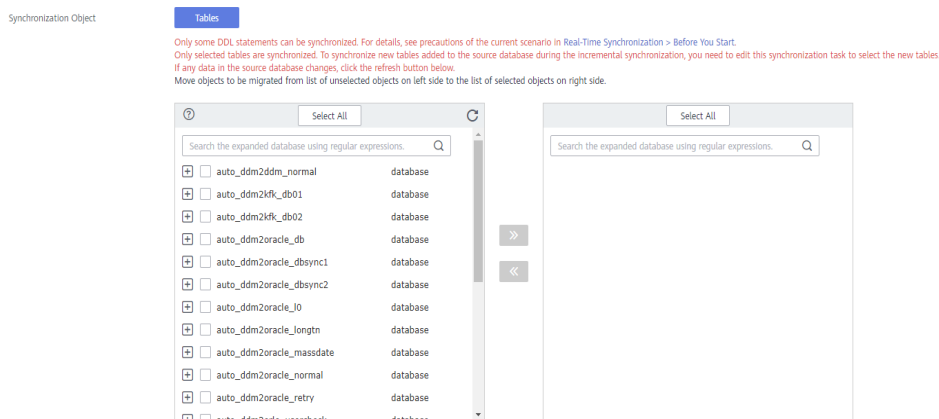



Table 4-61 Synchronization mode and object

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

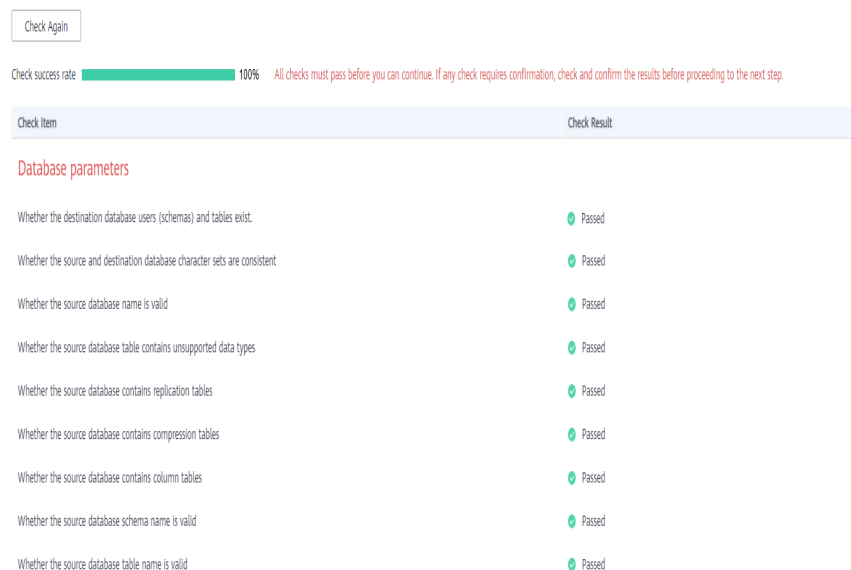
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-63 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-64 Task startup settings





Table 4-62 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.7 From DDM to Oracle

Supported Source and Destination Databases

Table 4-63 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 |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-64 Precautions

| Type | Restrictions |
|----------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. • The destination database user must have the following permissions: ALTER ANY INDEX, ALTER ANY TABLE, ALTER SESSION, ANALYZE ANY, COMMENT ANY TABLE, CREATE ANY INDEX, CREATE ANY TABLE, CREATE SESSION, DELETE ANY TABLE, DROP ANY TABLE, INSERT ANY TABLE, SELECT ANY TABLE, SELECT ANY DICTIONARY, SELECT ANY TRANSACTION, UPDATE ANY TABLE, and RESOURCE roles. |

| Type | Restrictions |
|------------------------|--|
| Synchronization object | <ul style="list-style-type: none"> ● The source database data can be synchronized. ● The source database table structure supports only full synchronization. ● Database objects other than table structures, indexes, and constraints cannot be synchronized. ● The destination Oracle database does not support empty strings, so the object to be synchronized cannot contain empty strings. ● The maximum number of columns supported by the source DDM database is 1017, while the maximum number of columns supported by the destination Oracle database is 1000. Therefore, the number of columns of objects to be synchronized cannot exceed 1000. |
| 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. |

| 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.• The destination database (account) must have the RESOURCE permission. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● If the data types are incompatible, the synchronization may fail. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The NOT NULL constraint of MySQL supports empty strings, while the NOT NULL constraint of Oracle does not. During a synchronization, if an empty string exists in a NOT NULL constraint field, delete the NOT NULL constraint from the destination Oracle database. ● If the source database contains a duplicate primary key, data synchronized to the destination database will be less than that in the source database. You must check and correct the data before starting the synchronization task. ● The varbinary, binary, and timestamp columns in the source table cannot contain primary key or unique constraints. ● If the length of the varchar data in the source database is greater than or equal to 667 characters, the varchar type will be converted to the clob type in the Oracle database. ● If 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. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • During an incremental synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During an incremental synchronization, if you need to modify the structure of the source table to be synchronized, you must modify the corresponding destination table structure. • During an incremental synchronization, do not perform the restoration operation on the source database. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-65 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 the warning, there are four input fields: "Region" (a dropdown menu), "Project" (a dropdown menu), "Task Name" (a text input field containing "DRS-8131"), and "Description" (a text area). 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." The form is enclosed in a dashed border.

Table 4-65 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 4-66 Synchronization instance details

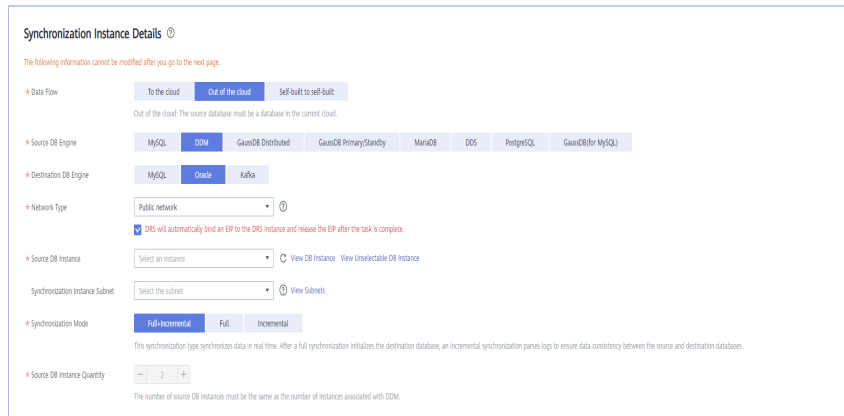


Table 4-66 Synchronization instance settings

| Parameter | Description |
|-----------------------|---|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select DDM . |
| Destination DB Engine | Select Oracle . |
| Network Type | Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The DDM instance you created. |

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |
| Source DB Instance Quantity | Specifies the number of DB instances bound to the source DDM database. |

- Task type

Figure 4-67 Task type



Table 4-67 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-68 Enterprise projects and tags




Table 4-68 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-69 Source database information

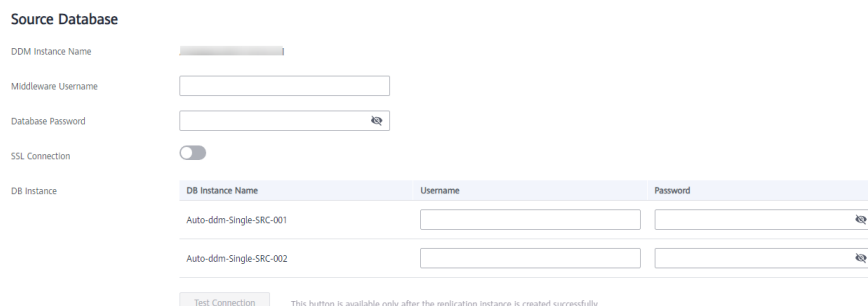


Table 4-69 Source database settings

| Parameter | Description |
|---------------------|---|
| DDM Instance Name | The DDM instance you selected when you create a synchronization task. The instance name cannot be changed. |
| Middleware Username | The username for accessing the source database. |
| Database Password | The database username and password are encrypted and stored in the system, and will be cleared after the task is deleted. |

| Parameter | Description |
|----------------|--|
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |
| DB Instance | The sharded database details. |

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 4-70 Destination database information

Destination Database


IP Address or Domain Name

For a RAC cluster, use a Scan IP address and specify Service Name to improve access performance.

Port

Database Service Name Service Name ▼ ?

Database Username

Database Password 

SSL Connection

This button is available only after the replication instance is created successfully.

Table 4-70 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | <p>The IP address or domain name of the destination database.</p> <p>NOTE</p> <p>For a RAC cluster, use a scan IP address to improve access performance.</p> |
| Port | The port of the destination database. Range: 1 - 65535 |

| Parameter | Description |
|-----------------------|--|
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-71 Synchronization mode

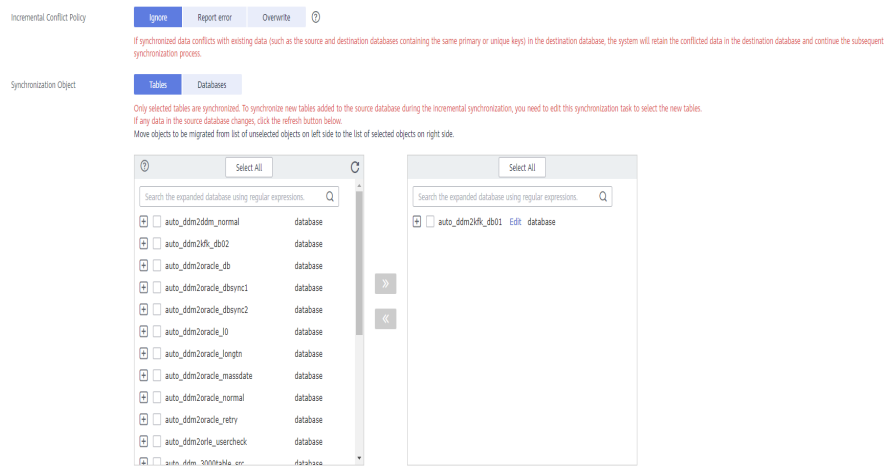




Table 4-71 Synchronization mode and object

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Data Processing** page, select the table object to be processed, enter the column name, type, and operation type to be added, confirm the information, and click **Next**. You can set related rules by referring to "Adding Synchronization Timestamp" in [Processing Data](#).

Figure 4-72 Processing data



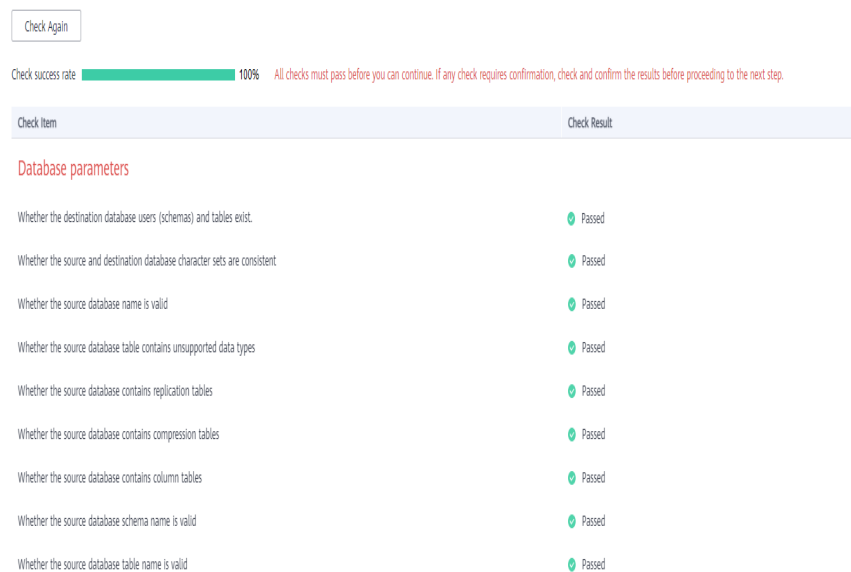
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-73 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 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-74 Task startup settings

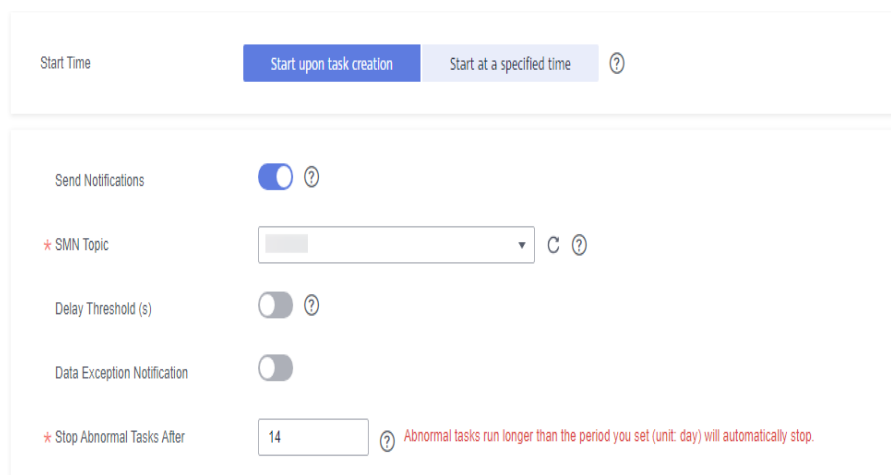



Table 4-72 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.8 From DDM to Kafka

Supported Source and Destination Databases

Table 4-73 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none"> • DDM instances | <ul style="list-style-type: none"> • Kafka |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-74 Precautions

| Type | Restrictions |
|------------------------|--|
| Database permissions | <ul style="list-style-type: none"> The source database DDM account must have at least one permission, for example, SELECT. The DDM physical sharded database account must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. |
| Synchronization object | <ul style="list-style-type: none"> The table data can be synchronized. |
| 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. You are advised to set auto.create.topics.enable of Kafka to false. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> • If the data types are incompatible, the synchronization may fail. • If the source database contains a duplicate primary key, data synchronized to the destination database will be less than that in the source database. Therefore, you must check and correct the data before starting the synchronization task. • After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be synchronized or the task fails. • During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During the synchronization, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. • DDL operations are supported during incremental synchronization. • During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. |

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-75 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. 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 downward arrow.
- Project:** A dropdown menu with a downward arrow.
- * Task Name:** A text input field containing "DRS-8131" and a help icon (question mark in a circle).
- Description:** A text area with a help icon (question mark in a circle).

At the bottom right of the form, there is a small number "0/256".

Table 4-75 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-76 Synchronization instance details




Table 4-76 Synchronization instance settings

| Parameter | Description |
|-----------------------|----------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select DDM . |
| Destination DB Engine | Select Kafka . |

| Parameter | Description |
|---------------------------------|--|
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | An available DDM instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. For details about the underlying working principles for incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p> |
| Source DB Instance Quantity | Specifies the number of DB instances bound to the source DDM database. Set this parameter based on the site requirements. |

- Task type

Figure 4-77 Task type



Table 4-77 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-78 Enterprise projects and tags

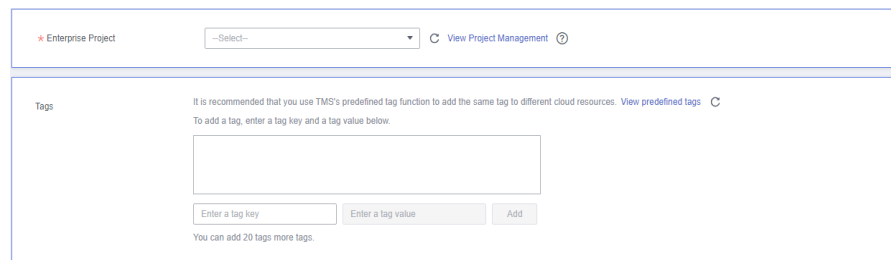


Table 4-78 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-79 Source database information

Source Database

DB Instance Name rds-000-source

Database Username

Database Password

✔ Test successful

Table 4-79 Source database settings

| Parameter | Description |
|-------------------|--|
| DB Instance Name | The DDM instance you selected when you create the synchronization task. The instance name cannot be changed. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

NOTE

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Figure 4-80 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

Table 4-80 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Security Protocol | Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication . |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 4-81 Synchronization mode

Topic Synchronization Policy: A specified topic

Topic: test

Synchronize Topic To: Partitions are identified by the hash values ...

Data Format in Kafka: Avro | JSON | **JSON-C**

Synchronization Object: **Tables** | Databases

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.

Select All

For tables, only expanded databases are searched.


- drstest_0000 database
- sbtest database
- test database

Select All

For tables, only expanded databases are searched.

Table 4-81 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 . |
| 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 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> |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table- and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

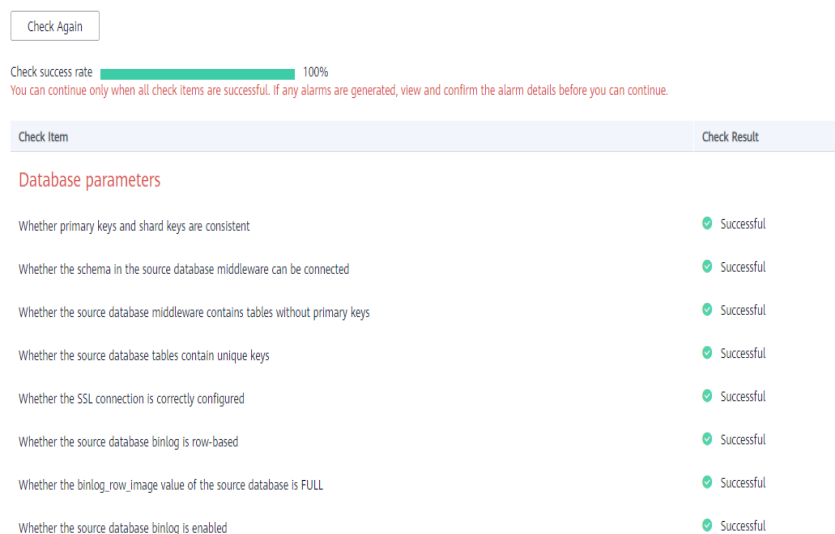
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-82 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-83 Task startup settings

Table 4-82 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.9 From DDS to MongoDB

Supported Source and Destination Databases

Table 4-83 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none">• DDS DB instances (versions 3.2, 3.4, 4.0, 4.2, and 4.4) | <ul style="list-style-type: none">• On-premises MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4)• ECS-hosted MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4)• Other cloud MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4) <p>NOTE The destination database version must be the same as or later than the source database version.</p> |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - Tables to be synchronized without a primary key may be locked for 3s.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

- If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-84 Precautions

| Type | Restrictions |
|----------------------|--|
| Database permissions | <p>Source database (minimum permissions):</p> <ul style="list-style-type: none"> • Permission requirements for incremental synchronization: <ul style="list-style-type: none"> - Replica set: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database. - Single node: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database. - Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. <p>Minimum permission requirements: The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database. If the destination database is a cluster instance, the database user must have the read permission for the config database.</p> |

| Type | Restrictions |
|------------------------|--|
| Synchronization object | <ul style="list-style-type: none"> ● Replica set: Only collections (including validator and capped and non-capped collections), indexes, and views can be synchronized. ● Cluster: Currently, only collections (including validator and capped and non-capped collections), shard keys, indexes, and views can be synchronized. ● Single node: Only collections (including validator and capped collections), indexes, and views can be synchronized. ● Collections that contain the <code>_id</code> field without indexes are not supported. ● The first parameter of <code>BinData()</code> cannot be 2. ● If ranged sharding is used, <code>maxKey</code> cannot be used as the primary key. ● If the source database is a cluster database, <code>DBPointer</code> and <code>DBRef</code> are not supported. ● Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the synchronization. |
| 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><code>\$</code> or spaces. The collection name and view name cannot start with <code>system.</code> or contain the dollar sign (<code>\$</code>). ● If the source database is a cluster, the balancer must be disabled and orphan documents must be deleted. For details, see How Do I Disable the Balancer? and How Do I Delete Orphaned Documents in MongoDB Sharded Clusters? |
| Destination database | <ul style="list-style-type: none"> ● Data cannot be synchronized from a newer version database to an older version database. ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● If the source is not a cluster instance, the following operations and commands are supported during incremental synchronization: <ul style="list-style-type: none"> - Creating and deleting databases - Adding, deleting, and updating documents - Creating and deleting collections - Creating and deleting indexes - Creating and deleting views - The <code>convertToCapped</code>, <code>collMod</code>, and <code>renameCollection</code> commands are supported. ● If the source is a cluster instance: <ul style="list-style-type: none"> - The cluster version cannot be earlier than 4.0. - Do not delete synchronization objects in the incremental synchronization phase. Otherwise, the synchronization task will fail. - The synchronization process consumes certain CPU and memory resources of the source database. Evaluate the source database resources in advance. - If the load on the source database is heavy, the processing speed of change streams cannot keep up with the oplog generation speed. As a result, DRS synchronization delay occurs. - Only the following DDLs are supported: drop database, drop collection and rename - In the incremental synchronization phase, the synchronization speed can reach up to 10,000 rows in a single table per second. ● During startup, the specified start point must be within the Oplog range. ● The destination database user must have the write permission. If the destination is a cluster instance, the database user must have the read permission for the config data. ● If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental synchronization, data consistency cannot be ensured when source and destination databases are in different time zone. ● The value of <code>block_compressor</code> is determined by <code>stats().wiredTiger.creationString.block_compressor</code> of the collection in the source database. If the destination database contains corresponding empty collections, the compression parameters will not be migrated. If the compression parameters in the source database are not supported by the destination database, configure the compression parameters based on <code>net.compression.compressors</code> of the destination |

| Type | Restrictions |
|------|--|
| | <p>database. If the storage engine of the destination database is not WiredTiger, DRS does not synchronize compression parameters.</p> <ul style="list-style-type: none"> ● If the destination is a replica set instance out of the cloud, enter information about all primary and secondary nodes to reduce the impact of a primary/secondary switchover on the synchronization task. If you enter information about primary and secondary nodes, ensure that all nodes belong to the same replica set instance. ● If the destination is a cluster instance out of the cloud, enter information about multiple mongos nodes to reduce the impact of the single-node failure on the synchronization task. In addition, multiple mongos nodes support load balancing. In addition, ensure that all mongos nodes belong to the same cluster instance. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● To ensure data consistency, do not modify the destination database (including but not limited to DDL and DML operations) during the entire synchronization process. ● During the synchronization, data rollback caused by a primary/standby switchover of the source database is not supported. ● The system database cannot be synchronized. The username and role must be manually created in the destination database. ● Documents larger than 16 MB in the source database cannot be inserted or updated during incremental synchronization. ● In the incremental synchronization phase, concurrent replay is performed at the collection level to maintain the synchronization performance. In the following scenarios, only single-thread write is supported and concurrent replay is not supported: <ul style="list-style-type: none"> - The collection index contains a unique key. - The value of capped of the collection attribute is true. <p>In either of the preceding scenarios, the task delay may increase.</p> ● To prevent loopback, the to-the-cloud migration task and out-of-cloud synchronization task cannot be started at the same time. ● During row comparison, if an orphan document exists in a cluster instance or chunks are being synchronized, the number of returned rows is incorrect and the comparison results are inconsistent. For details, see MongoDB official documentation. |

Procedure

This section uses real-time incremental synchronization from DDS to MongoDB as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-84 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 text area with a height of 0/256 characters and a help icon.

Table 4-85 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-85 Synchronization instance details




Table 4-86 Synchronization instance settings

| Parameter | Description |
|---------------------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select DDS. |
| Destination DB Engine | Select MongoDB . |
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | An available DDS instance |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. For details about the underlying working principles for incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p> |

- Task type

Figure 4-86 Task type



Table 4-87 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-87 Enterprise projects and tags

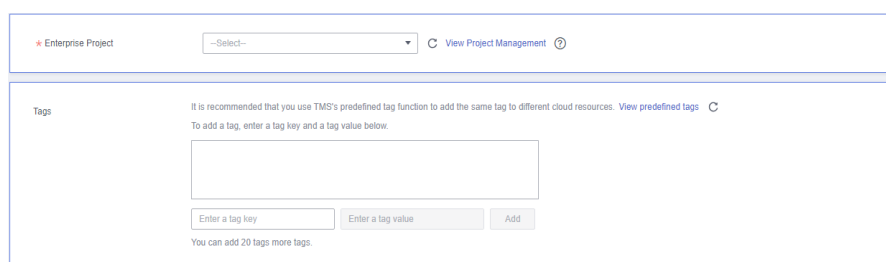


Table 4-88 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-88 Source database information

Source Database

DB Instance Name

Authentication Database

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 4-89 Source database settings

| Parameter | Description |
|-------------------------|--|
| DB Instance Name | The DDS instance you selected when creating the task. This parameter cannot be changed. |
| Authentication Database | The name of the authentication database. For example: The default authentication database of Huawei Cloud DDS instance is admin . |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.

Figure 4-89 Destination database information

The screenshot shows a configuration form for a destination database. It contains the following elements:

- IP Address or Domain Name:** A text input field with a help icon. Below it is a red note: "Ensure that the entered addresses belong to the same DB instance."
- Authentication Database:** A text input field.
- Database Username:** A text input field.
- Database Password:** A text input field with a visibility toggle icon.
- SSL Connection:** A blue toggle switch that is currently turned on. Below it is a red note: "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:** A text input field followed by a "Select" button.
- Test Connection:** A button at the bottom of the form.

Table 4-90 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | <p>IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535</p> <p>You can enter up to 3 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080. Ensure that the entered IP addresses or domain names belong to the same instance.</p> <p>NOTE If multiple IP addresses or domain names are entered, the test connection is successful as long as one IP address or domain name is accessible. Therefore, you must ensure that the IP address or domain name is correct.</p> |
| Authentication Database | The name of the destination database. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Task** page, select the start point and synchronization objects, and click **Next**.

Figure 4-90 Synchronization mode

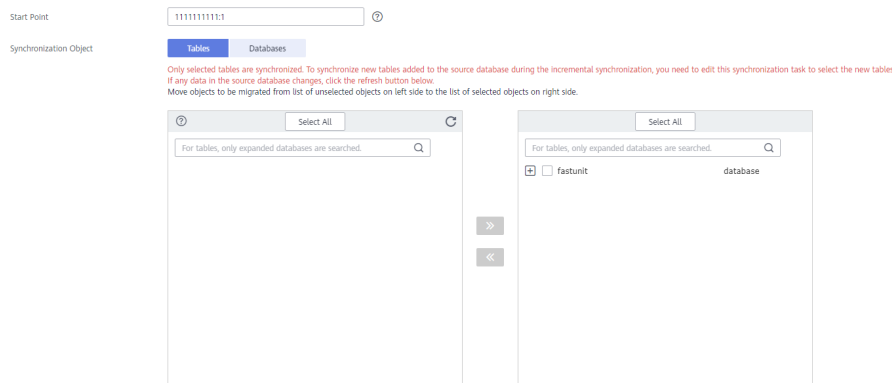



Table 4-91 Synchronization Object

| 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 <code>db.getSiblingDB("local").oplog.rs.find()</code> 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-91 Pre-check

Check success rate 100% All checks must pass before you can continue. If any check requires confirmation, check and confirm the results before proceeding to the next step.

| Check Item | Check Result |
|--|--------------|
| Destination database storage space | |
| Whether the destination database has sufficient storage space | ✔ Passed |
| Conflict | |
| Whether collections in both the source and destination databases are not capped | ✔ Passed |
| Whether the destination database contains a non-empty collection with the same name as that in the source database | ✔ Passed |
| Whether the same view names exist in both the source and destination databases | ✔ Passed |
| Object dependency | |
| Whether the source database referenced roles pass the check | ✔ Passed |
| Whether the source database referenced accounts pass the check | ✔ 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 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-92 Task startup settings

Start Time

?

Send Notifications

?

* SMN Topic

C ?

Delay Threshold (s)

?

Data Exception Notification


* Stop Abnormal Tasks After

?
Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 4-92 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.10 From DDS to Kafka

Supported Source and Destination Databases

Table 4-93 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none"> • DDS DB instances (versions 4.0, 4.2, and 4.4) | <ul style="list-style-type: none"> • Kafka |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-94 Precautions

| Type | Restrictions |
|----------------------|---|
| Database permissions | Source database (minimum permissions): <ul style="list-style-type: none"> • Permission requirements for incremental synchronization: The source database user must have the readAnyDatabase permission for the admin database. |
| Source database | <ul style="list-style-type: none"> • During the incremental synchronization, the Oplog of the source database must be enabled. • If the storage space is sufficient, store the source database Oplog for as long as possible. The recommended retention period is three days. • The source database name cannot contain \\. "\$ or spaces. The collection name cannot start with system. or contain the dollar sign (\$). • If the source database is a cluster, the balancer must be disabled and orphan documents must be deleted. For details, see How Do I Disable the Balancer? and How Do I Delete Orphaned Documents in MongoDB Sharded Clusters? |
| Destination database | <ul style="list-style-type: none"> • The destination database is a Kafka database. • You are advised to set auto.create.topics.enable of Kafka to false. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> • The synchronization process consumes certain CPU and memory resources of the source database. Evaluate the source database resources in advance. • If the source database is under a heavy load, the processing speed of change streams cannot keep up with the oplog generation speed. As a result, DRS synchronization delay occurs. • The following DMLs are supported: insert, update, replace, and delete. • Only the following DDLs are supported: drop database, drop collection and rename • In the incremental synchronization phase, the synchronization speed can reach up to 10,000 rows in a single table per second. • During startup, the specified start point must be within the Oplog range. • During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • To ensure data consistency, do not modify the destination database (including but not limited to DDL and DML operations) during the entire synchronization process. • Documents larger than 16 MB in the source database cannot be inserted or updated during incremental synchronization. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. • During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. |

Procedure

This section uses real-time incremental synchronization from DDS to Kafka as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-93 Synchronization task information

Table 4-95 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 4-94 Synchronization instance information

Table 4-96 Synchronization instance settings

| Parameter | Description |
|------------------|----------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select DDS . |

| Parameter | Description |
|---------------------------------|--|
| Destination DB Engine | Select Kafka . |
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | An available DDS instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. For details about the underlying working principles for incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p> |

- Task type

Figure 4-95 Task type



Table 4-97 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-96 Enterprise projects and tags

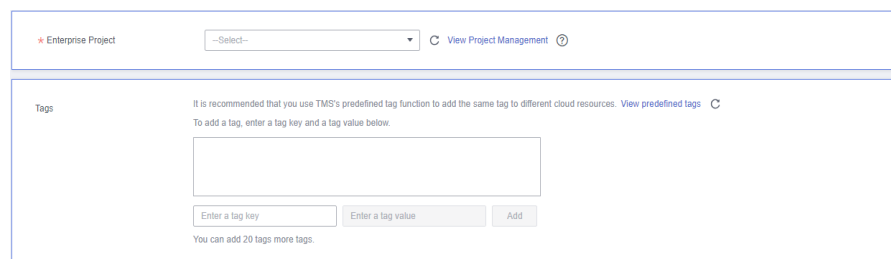


Table 4-98 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

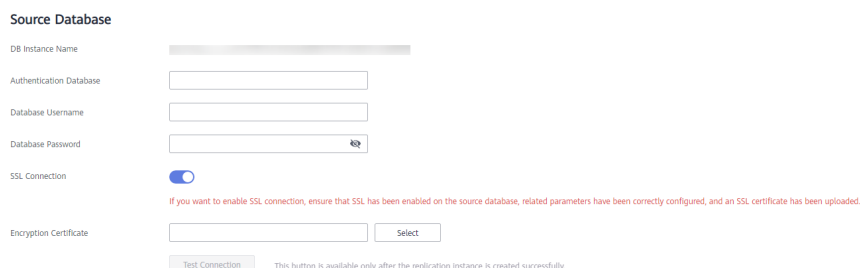
| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-97 Source database information



Source Database

DB Instance Name

Authentication Database

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 4-99 Source database settings

| Parameter | Description |
|-------------------------|--|
| DB Instance Name | The DDS instance you selected when creating the task. This parameter cannot be changed. |
| Authentication Database | The name of the authentication database. For example: The default authentication database of Huawei Cloud DDS instance is admin . |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |


| Parameter | Description |
|----------------|--|
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The IP address, 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 4-98 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.


Security Protocol 

Table 4-100 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Security Protocol | Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication . |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Task** page, select the start point and synchronization objects, and click **Next**.

Figure 4-99 Synchronization Mode

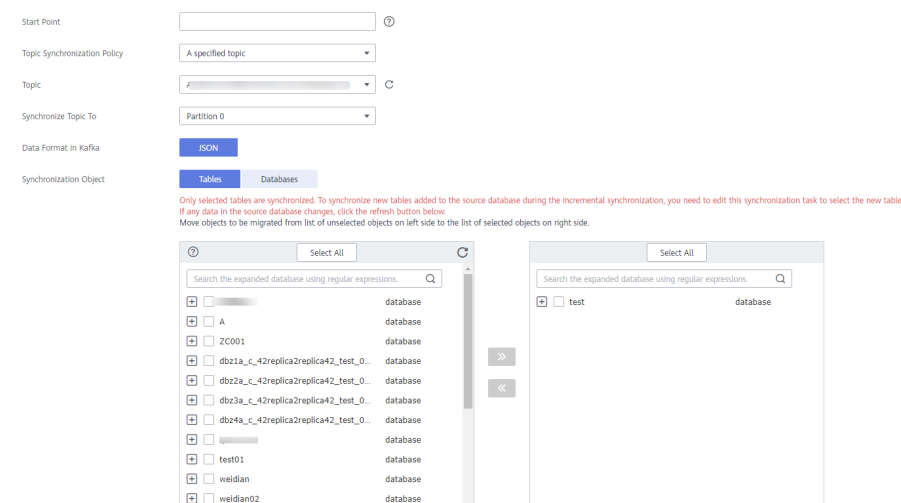



Table 4-101 Synchronization Object

| 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 <code>db.getSiblingDB("local").oplog.rs.find()</code> to query the source database oplog. The format of the <code>ts</code> 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. |
| Topic Synchronization Policy | Topic synchronization policy. You can select A specific topic . |
| Topic | Select the topic to be synchronized to the destination database. |
| Synchronize Topic To | <p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic. |
| Data Format in Kafka | <p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

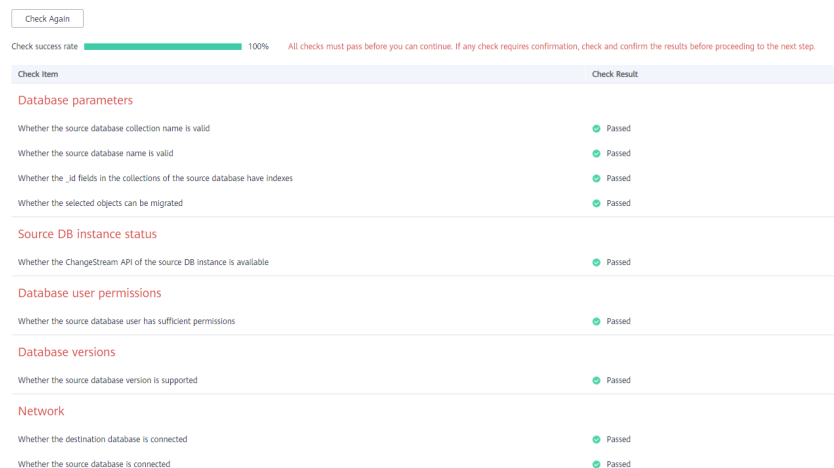
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-100 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-101 Task startup settings





Table 4-102 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.11 From PostgreSQL to PostgreSQL

Supported Source and Destination Databases

Table 4-103 Supported databases

| Source DB | Destination DB |
|---|---|
| RDS for PostgreSQL (versions 9.5, 9.6, 10, 11, 12, 13, 14 and 15) | <ul style="list-style-type: none"> On-premises databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15) ECS-hosted databases (PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14 and 15) <p>NOTE The major version of the destination database must be the same as or later than that of the source database.</p> |

Supported Synchronization Objects

Table 4-104 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-104 Supported synchronization objects

| Type | Notes |
|---------|--|
| Objects | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. Multiple DRS tasks are required to synchronize multiple databases. ● Supported field types: Digit, currency, character, binary, date/time, boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range. ● Scope of full synchronization <ul style="list-style-type: none"> - The following objects are supported in the database-level synchronization: schemas, tables, indexes, constraints, views, materialized views, sequences, stored procedures, rules, triggers, foreign keys, sorting rules, plug-ins, code conversion information, aggregate functions, operators, statistics extension, conversion information, text search configurations, functions, data types, type conversion, event triggers, text search parsers, and text search templates During the table-level synchronization, only tables, views, materialized views, sequences, and common indexes can be synchronized. During object file import, tables can be synchronized. - System objects such as system schemas, system tables, system users, and system functions cannot be synchronized. Other objects such as tablespaces, foreign data wrappers, foreign servers, user mappings, releases, subscriptions, and users cannot be synchronized. Object permissions cannot be synchronized. <p>NOTE The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none"> ● Object name: The database name cannot contain "+" %?\<>, the schema name and table name cannot contain ".'\<>, and the column name cannot contain double quotation marks (") and single quotation marks ('). ● Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers. ● Function: Do not synchronize C language functions or functions with the leakproof or support attribute. ● Plug-in: The metadata of plug-ins is not synchronized. ● Data type: Basic data types are not synchronized. ● Type conversion: The binary coercion type cannot be converted. <ul style="list-style-type: none"> ● Scope of incremental synchronization <ul style="list-style-type: none"> - Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized. - Some DDL statements can be synchronized, including TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA, CREATE TABLE, DROP TABLE, ALTER TABLE |

| Type | Notes |
|------|---|
| | <p>(including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, ALTER VIEW, COMMENT ON COLUMN, COMMENT ON TABLE, COMMENT ON SCHEMA, COMMENT ON SEQUENCE, COMMENT ON INDEX, and COMMENT ON VIEW.</p> <p>During table-level synchronization, only the following DDL operations can be synchronized: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, COMMENT ON COLUMN, COMMENT ON TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME).</p> <ul style="list-style-type: none"> - Not synchronized: DML statements of unlogged tables and temporary tables <p>NOTE The source database captures DDL statements using event triggers and records them in specific tables, so you need to create event triggers and functions in the source database in advance. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</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 4-105](#). 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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-105 Database account permission

| Type | Full | Full+Incremental |
|---------------------------|--|--|
| Source database user | Database CONNECT permission, schema USAGE permission, table SELECT permission, and sequence SELECT permission | <p>The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permissions for sequences, the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections</p> <p>NOTE The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration.</p> |
| Destination database user | <ul style="list-style-type: none"> ● Database-level: <ul style="list-style-type: none"> - If the destination database is not PostgreSQL, the CREATEDB permission is required. - If the destination database is PostgreSQL, the CONNECT and CREATE permissions on PostgreSQL databases and the USAGE and CREATE permissions on public schemas are required. ● Table-level: <ul style="list-style-type: none"> - To synchronize databases, the CREATEDB permission is required. - To synchronize a schema, the CONNECT and CREATE permissions for the database that contains the schema are required. - To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the object are required. <p>NOTE If the destination database version is PostgreSQL 15 and no database is created on the destination database for database-level synchronization and table-level synchronization, full or incremental synchronization may fail due to lack of the USAGE and CREATE permissions of the public schema.</p> | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 4-106 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The partition table trigger of the source database cannot be set to disable. - For a full synchronization task, the source database can be a standby database, but hot_standby_feedback must be set to on. For an incremental synchronization task, the source database cannot be a standby database. - To perform incremental synchronization: The wal_level value of the source database must be logical. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Source database object requirements: <ul style="list-style-type: none"> - Triggers with the same name cannot exist in the source database. - The objects that have dependencies must be synchronized at the same time. Otherwise, the synchronization may fail. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The block_size value of the destination database must be greater than or equal to that of the source database. - The lc_monetary values of the source and destination databases must be the same. - To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to the original one. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination database cannot contain objects with the same type and name as the objects to be synchronized, including databases, schemas, and tables. System databases, system schemas, and system tables are excluded. - The destination table can contain more columns than the source table. However, the following failures must be avoided: |

| Type | Constraints |
|------|--|
| | <p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value +original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_+ hash value + original constraint name (which may be truncated) + _key. - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - For a full+incremental or incremental synchronization task, if an internal error occurs during the pre-check and the task stops before it is started, check and delete the streaming replication slot by referring to Forcibly Stopping Synchronization of PostgreSQL to avoid residual streaming replication slots in the source database. - After a full+incremental or incremental task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - If you choose to synchronize DDL statements, ensure that the DDL statements executed on the source database are compatible with the destination database. <p>NOTE</p> <p>DDL statements are captured using event triggers in the source database, recorded in a specific table, and then synchronized to the destination database. You need to create event triggers and functions in the source database before starting a task. For details, see Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</p> |

| 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 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 synchronization task is complete, the sequence values are compared based on logical consistency. That is, if an auto-increment sequence value is used only when the sequence value of the destination database is greater than or equal to that of the source database, or an auto-decrement sequence value is used only when the sequence value of the destination database is less than or equal to that of the source database, the sequence values in the source database are consistent with those in the destination database. - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. - If the value of session_replication_role of the destination database is replica when the full+incremental synchronization task is complete, change the value to the original one. ● Forcibly stop a task: <ul style="list-style-type: none"> - You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - If the value of session_replication_role of the destination database is replica, change it to the original value to forcibly stop the full+incremental synchronization task. - The naming rule of a logic replication slot is <code>drs_unique_ID</code>. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |
| Troubleshooting | <ul style="list-style-type: none"> ● If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting. |

Procedure

This section uses out-of-cloud synchronization from PostgreSQL to PostgreSQL as an example to describe how to configure a real-time synchronization task in the VPC network scenario.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-102 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 4-107 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| 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 4-103 Synchronization instance details




Table 4-108 Synchronization instance settings

| Parameter | Description |
|---------------------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select PostgreSQL . |
| Destination DB Engine | Select PostgreSQL . |
| Network Type | Available options: VPC, Public network, and VPN or Direct Connect . VPC is used as an example. <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source 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 Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-104 Task type



Table 4-109 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-105 Enterprise projects and tags




Table 4-110 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.


- Source database configuration

Figure 4-106 Source database information

Source Database

DB Instance Name pg-12-for-autotest ([redacted])

Database Username

Database Password 

✔ Test successful

Table 4-111 Source database settings

| Parameter | Description |
|-------------------|---|
| DB Instance Name | The RDS for PostgreSQL instance you selected when creating the task. The parameter cannot be changed. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

- Destination database configuration

Figure 4-107 Destination database information

Destination Database


VPC [View VPC](#)

Subnet [View Subnets](#)

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 4-112 Destination database settings

| Parameter | Description |
|---------------------------|---|
| VPC | A dedicated virtual network in which the destination database is located. It isolates networks for different services. You can select an existing VPC or create a VPC. |
| Subnet | A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the destination database resides. You need to enable DHCP for creating the destination database subnet. |
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. |

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-108 Synchronization Mode

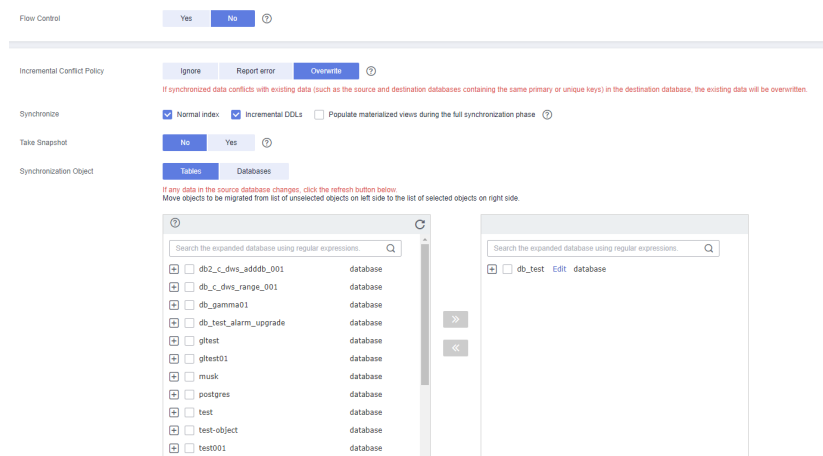
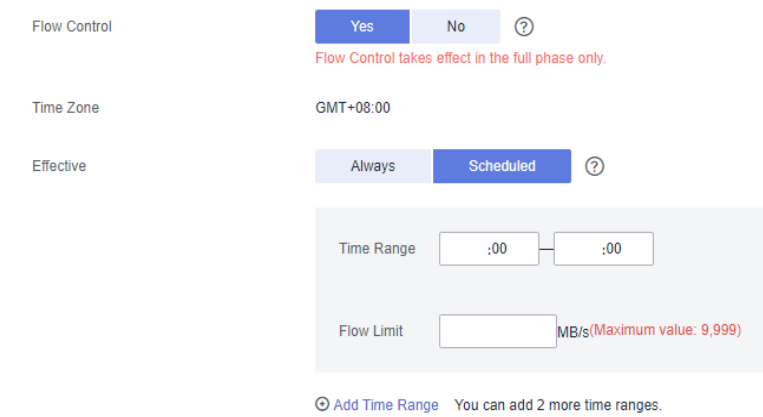



Table 4-113 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-109 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 |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| Synchronize | <p>Available options: 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

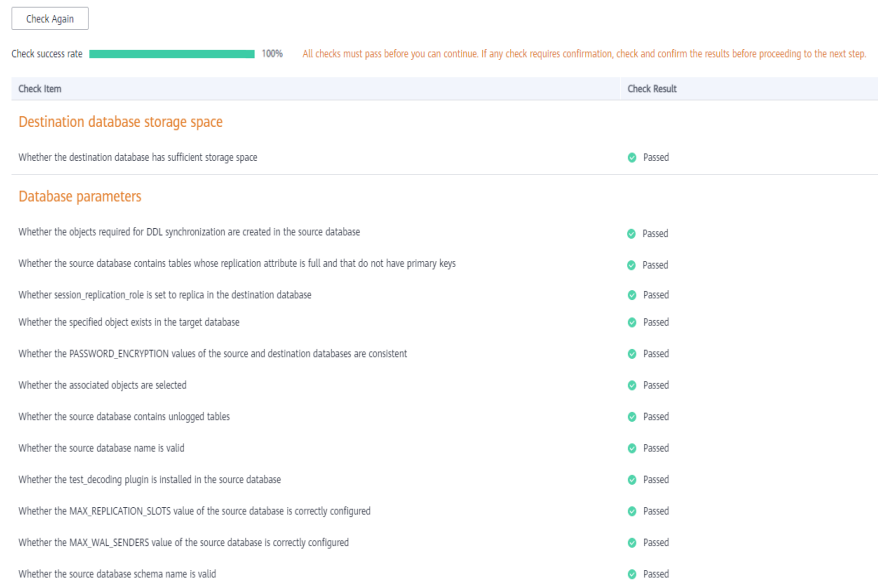
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-110 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-111 Task startup settings

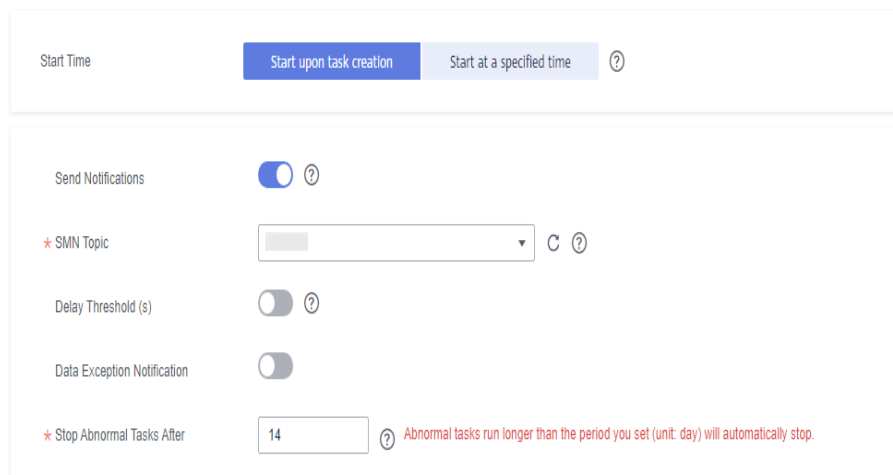



Table 4-114 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.12 From PostgreSQL to Kafka

Supported Source and Destination Databases

Table 4-115 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 4-116](#) 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 | 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 4-117](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-117 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 4-118 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> The wal_level value of the source database must be logical. The test_decoding plug-in has been installed on the source database. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - The source must be an RDS for PostgreSQL instance. - During database-level synchronization, incremental synchronization tasks cannot be edited. - During table-level synchronization, incremental synchronization tasks can be edited, but the database cannot be changed. - Before starting a synchronization task, ensure that no long transaction is started in the source database. Otherwise, the creation of the logical replication slot will be blocked, leading the task to fail. - After a task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. |

| Type | Constraints |
|-----------------------------|---|
| Incremental synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not delete the primary key of the source database table. Otherwise, incremental data may be lost or the task may fail. ● Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be lost or the task may fail. ● Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, incremental data may be lost or the task may fail. ● During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail. ● If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. |
| Stopping a task | <ul style="list-style-type: none"> ● Stop a task normally: <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. ● Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.

- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-112 Synchronization task information

Table 4-119 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-113 Synchronization instance details

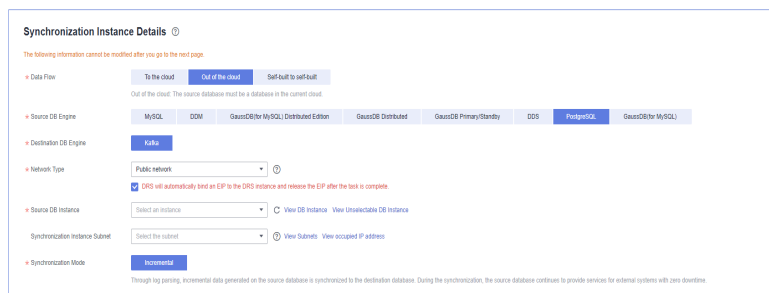


Table 4-120 Synchronization instance settings

| Parameter | Description |
|---------------------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select PostgreSQL . |
| Destination DB Engine | Select Kafka . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | Specifies an RDS PostgreSQL DB instance you have created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime. |

- Task type

Figure 4-114 Task type




Table 4-121 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-115 Enterprise projects and tags




Table 4-122 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 4-116 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 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 4-117 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 4-118 Synchronization mode

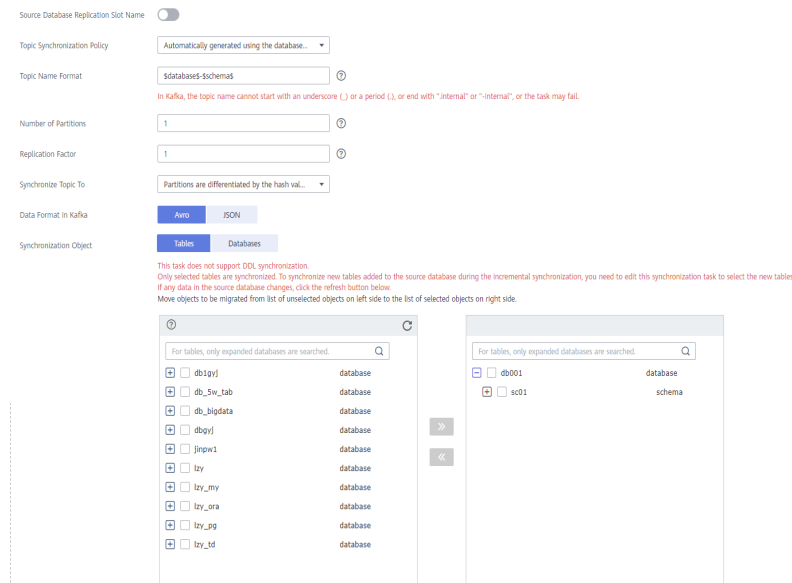



Table 4-125 Synchronization object

| 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 . |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

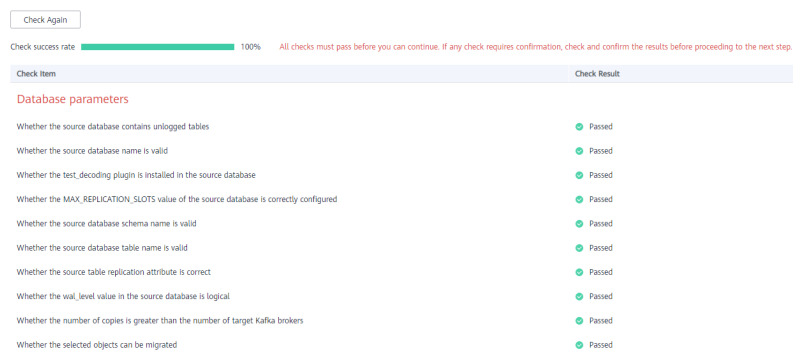
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-119 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-120 Task startup settings

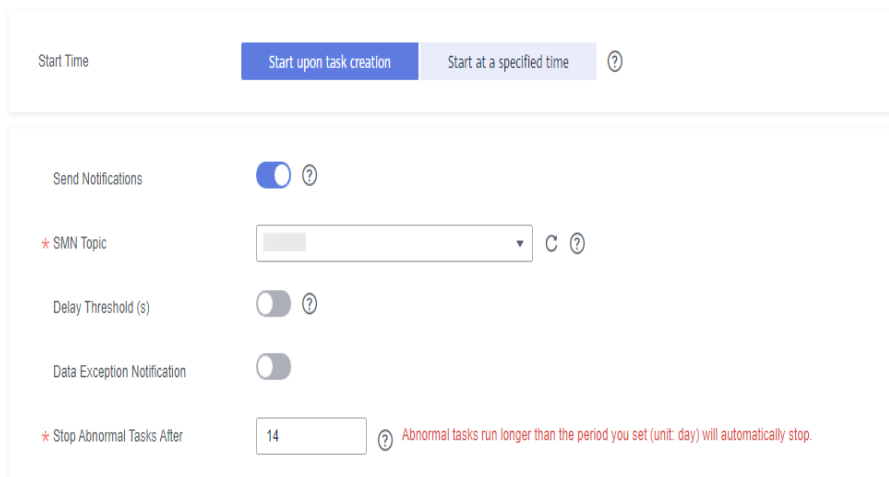



Table 4-126 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.13 From GaussDB Distributed to MySQL

Supported Source and Destination Databases

Table 4-127 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 4-128 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 | 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. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 4-129](#). 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. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user must have the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|--|--|
| | | details, see Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-130 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\`' \?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the 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 |

| Type | Restrictions |
|----------------------|--|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● 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 names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| Full synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. |

| 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 GaussDB Developer Guide. |
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Do not limit the synchronization speed during data comparison. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. |

| Type | Restrictions |
|-----------------|--|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-121 Synchronization task information

Table 4-131 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-122 Synchronization instance details

Table 4-132 Synchronization instance settings

| Parameter | Description |
|------------------|-------------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Distributed . |

| Parameter | Description |
|---------------------------------|---|
| Destination DB Engine | Select MySQL . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | The GaussDB distributed instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Incremental, and Full</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-123 Task type



Table 4-133 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-124 Enterprise projects and tags




Table 4-134 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-125 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-135 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 4-126 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 4-136 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-127 Synchronization mode

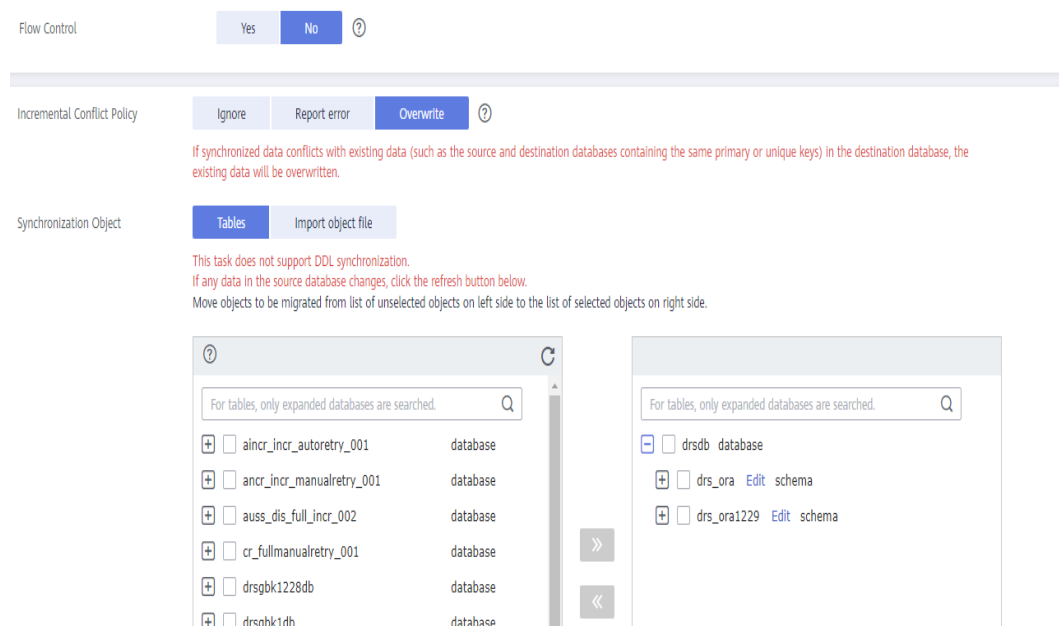
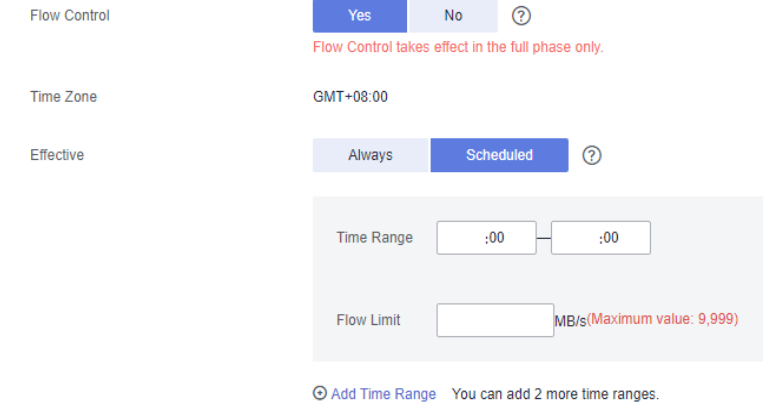



Table 4-137 Synchronization Object

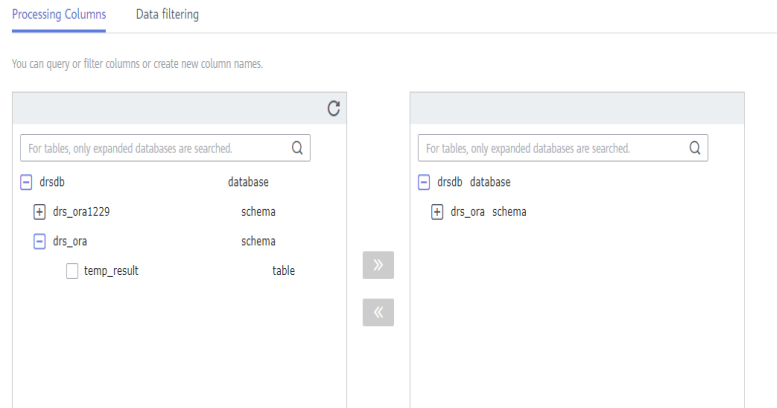
| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-128 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 |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-129 Processing data



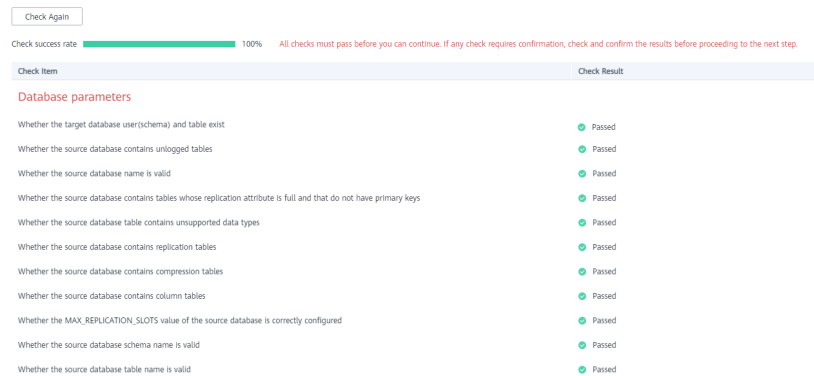
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-130 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-131 Task startup settings

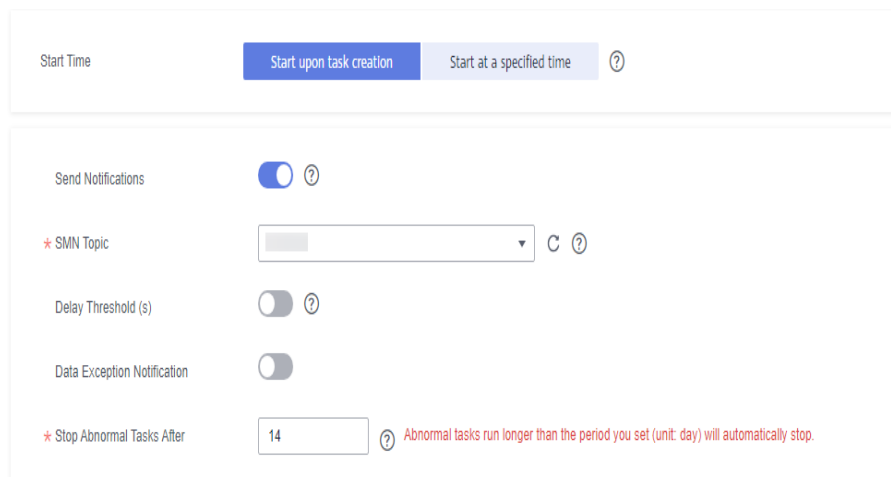



Table 4-138 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.14 From GaussDB Distributed to Oracle

Supported Source and Destination Databases

Table 4-139 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 4-140](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-140 Supported synchronization objects

| Type | Synchronization Scope |
|-----------------------|--|
| Synchronization scope | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported 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. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \\?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - 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 4-141](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source 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 4-141 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|---|--|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role. | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 4-142 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. - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails. - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task |

| Type | Restrictions |
|----------------------|--|
| | <p>failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs.</p> <ul style="list-style-type: none"> - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. • For tables containing large fields (including blob, clob, nclob, long, and long raw), large fields are not used as where conditions for incremental data. You are advised to ensure that all fields except large fields have the unique constraint. • If a table in the destination database contains both the lob and long types or both the lob and long raw types, and the length of the long or long raw type to be inserted or updated is greater than 4000, the insertion or update fails. You are advised to use the clob or blob type. |
| Data processing | <ul style="list-style-type: none"> • During column processing, the primary key, unique key, and distribution column cannot be filtered out. |
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison. |

| Type | Restrictions |
|-----------------|--|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-132 Synchronization task information

Table 4-143 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-133 Synchronization instance details

Table 4-144 Synchronization instance settings

| Parameter | Description |
|------------------|-------------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Distributed . |

| Parameter | Description |
|---------------------------------|---|
| Destination DB Engine | Select Oracle . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | The GaussDB distributed instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| 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. |

- Task type

Figure 4-134 Task type



Table 4-145 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-135 Enterprise projects and tags




Table 4-146 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-136 Source database information

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 4-137 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 4-148 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-138 Synchronization mode

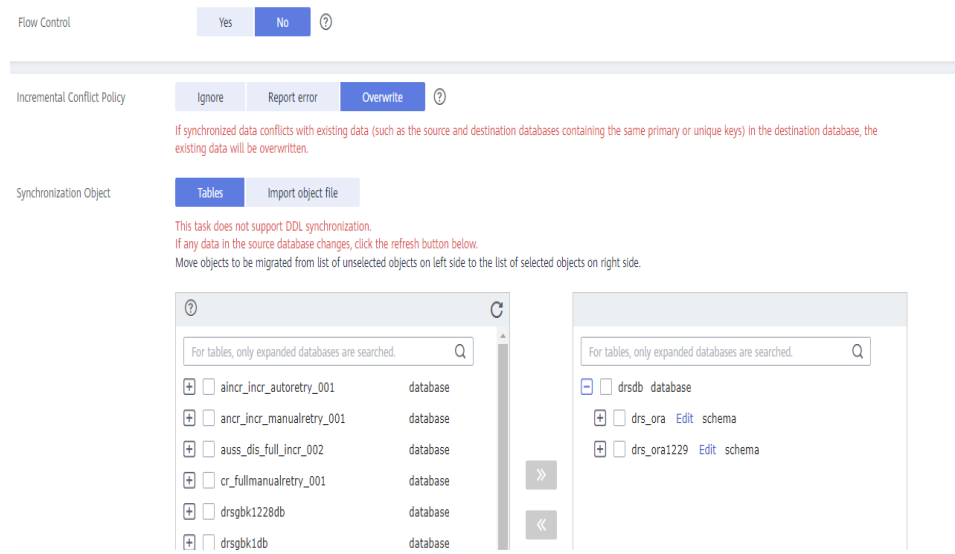
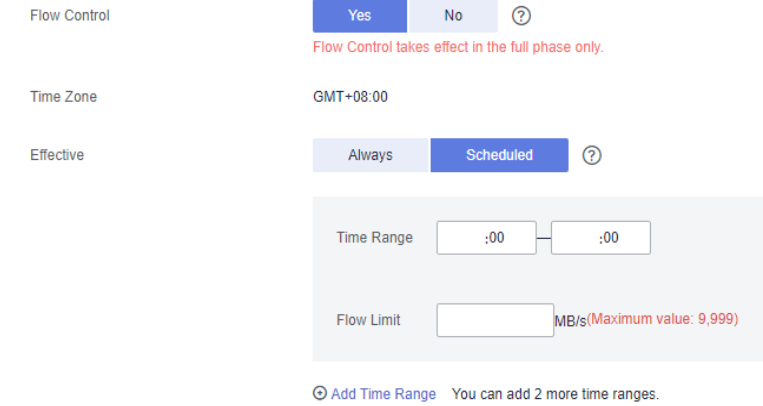



Table 4-149 Synchronization Object


| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-139 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-140 Processing data



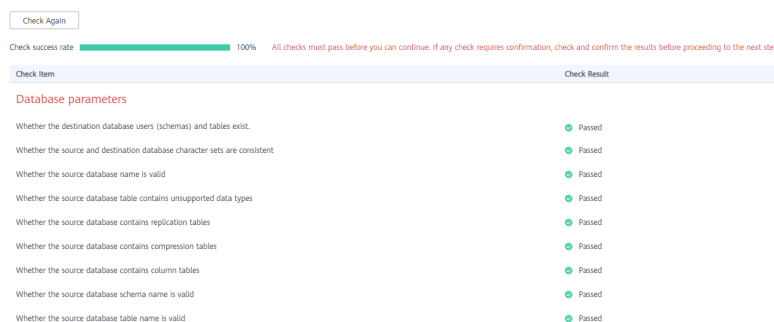
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-141 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-142 Task startup settings





Table 4-150 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.15 From GaussDB Distributed to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-151 Supported databases

| Source DB | Destination DB |
|---------------------|----------------------|
| GaussDB distributed | GaussDB(DWS) cluster |

Supported Synchronization Objects

[Table 4-152](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-152 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, temporary tables, and level-2 partitioned tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !/?!. The column name |

| Type | Synchronization Scope |
|------|---|
| | <p>cannot contain double quotation marks ("), single quotation marks ('), or periods (.).</p> <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-153](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source 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 4-153 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, and SELECT or UPDATE (required only for tables without primary keys) permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|--|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | The user must have the sysadmin role or the following permissions: <ul style="list-style-type: none"> The CONNECT and CREATE permissions on databases The USAGE permission on schemas The INSERT, DELETE, UPDATE, and SELECT permissions on tables | | |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-154 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<>\\` ?!. - The distribution key of the source database cannot contain commas (,) and spaces. - The primary key and unique constraint of the source database table must contain the distribution column. Otherwise, the table structure cannot be synchronized. You need to manually create the table structure in the destination database and retry the task. - Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario and causes value comparison to be unavailable. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - Before executing an incremental synchronization task, create schemas and tables in the destination database. - For an incremental synchronization task, disable foreign keys for tables in the destination database to ensure that the sequence in which data is written to the destination database is the same as that in the source database. Otherwise, the synchronization task may fail. - The destination table can contain more columns than the source table. However, the following failures must be avoided: |

| Type | Constraints |
|----------------------|--|
| | <p>Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

| Type | Constraints |
|-----------------|---|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num() command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-143 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, 0/256]

Table 4-155 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-144 Synchronization instance details

Synchronization Instance Details

Data Flow: To the cloud | **Out of the cloud** | Self-built self-built

Source DB Engine: MySQL | ODM | GaussDB for MySQL Distributed Edition | **GaussDB Distributed** | GaussDB Primary/Standby | DDS | PostgreSQL | GaussDB for MySQL

Destination DB Engine: MySQL | Oracle | **GaussDB Distributed** | GaussDB Distributed | Kafka

Network Type: Public network

Source DB Instance: No DB instance available | View DB Instance | View Unselectable DB Instance

Synchronization Instance Subnet: Select the subnet | View Subnets

Synchronization Mode: Full | **Full** | Incremental

Table 4-156 Synchronization instance settings

| Parameter | Description |
|------------------|-------------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Distributed . |

| Parameter | Description |
|---------------------------------|---|
| Destination DB Engine | Select GaussDB(DWS) . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for data synchronization between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The GaussDB distributed instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-145 Task type



Table 4-157 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-146 Enterprise projects and tags




Table 4-158 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-147 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-159 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 4-148 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 4-160 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). 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 4-149 Synchronization mode

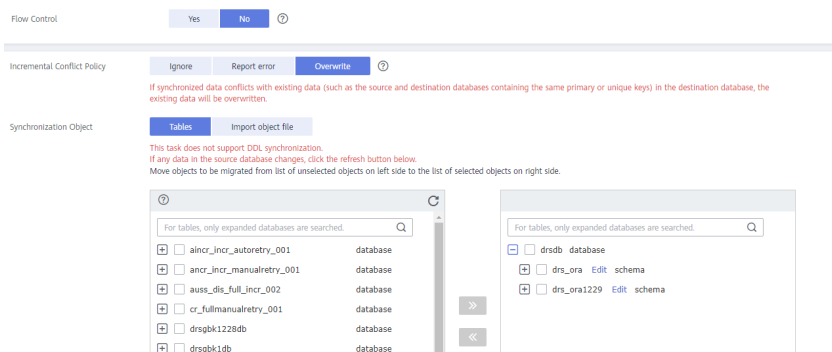
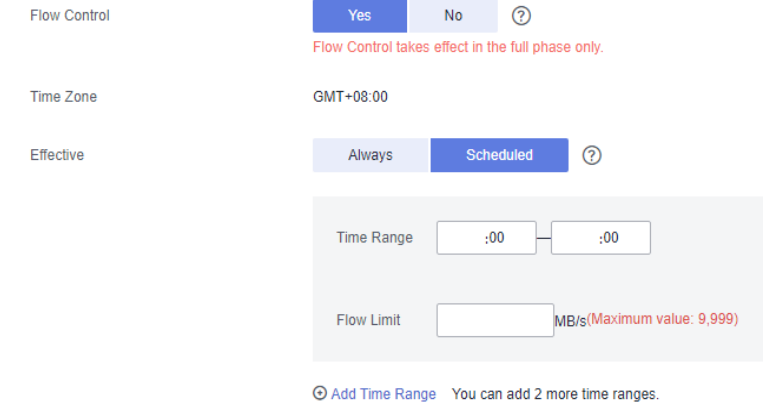



Table 4-161 Synchronization Object

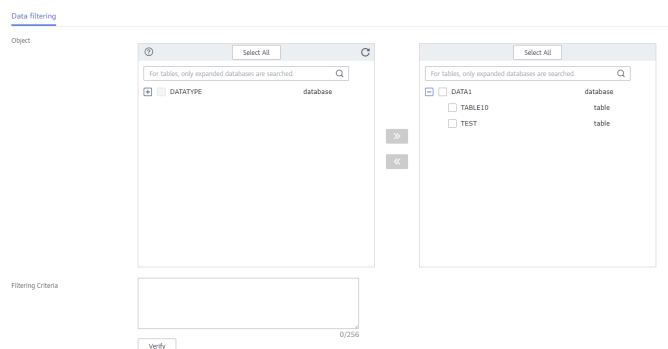
| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-150 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-151 Processing data



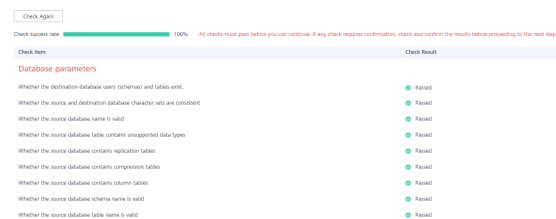
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-152 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-153 Task startup settings





Table 4-162 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.16 From GaussDB Distributed to Kafka

Supported Source and Destination Databases

Table 4-163 Supported databases

| Source DB | Destination DB |
|---------------------|---------------------|
| GaussDB distributed | Kafka 0.11 or later |

Supported Synchronization Objects

[Table 4-164](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-164 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 • The supported fields are BIGINT, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE and TINYINT. • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of the selected table can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-165](#).

DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source 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 4-165 Database user permission

| Type | Incremental |
|----------------------|---|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. • The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. • Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 4-166 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\` \?!. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|--|
| Incremental synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During migration of table-level objects, you are not advised to rename the tables. • Replication of interval partition tables is not supported. • The name of a primary key column cannot be changed. • After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. • Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-154 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, 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 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." The Task Name field has a character count icon showing "0/50". The Description field has a character count icon showing "0/256".

Table 4-167 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-155 Synchronization instance details

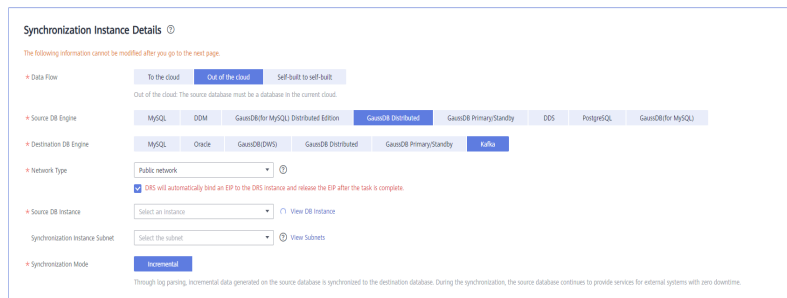


Table 4-168 Synchronization instance settings

| Parameter | Description |
|-----------------------|---|
| Data Flow | Select Out of the cloud . |
| 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 <ul style="list-style-type: none"> – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | The GaussDB distributed instance you created. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Synchronization Mode | <ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime. |

- Task type

Figure 4-156 Task type



Table 4-169 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-157 Enterprise projects and tags

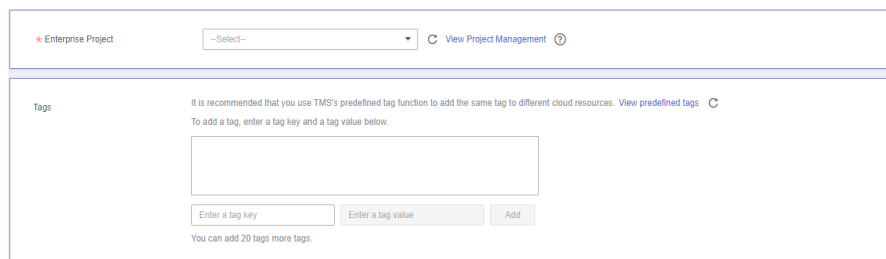


Table 4-170 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-158 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 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 4-159 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 4-172 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). 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 4-160 Synchronization mode

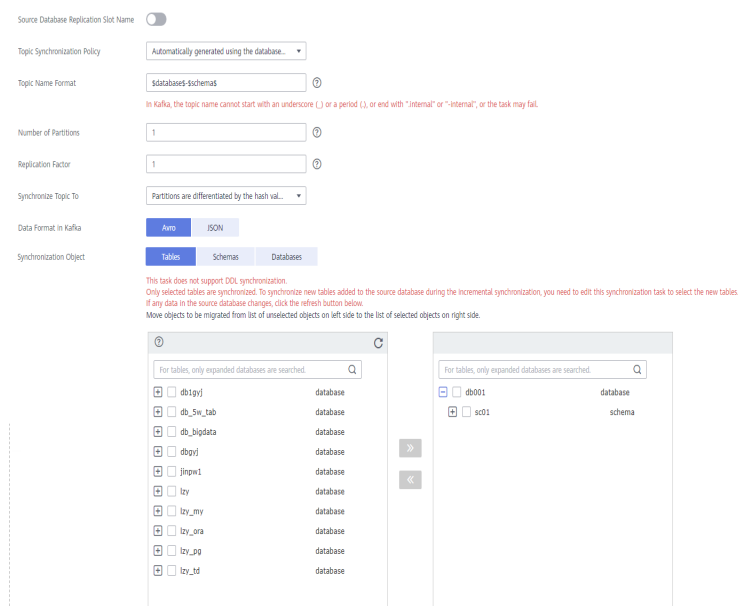



Table 4-173 Synchronization Object

| 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. |

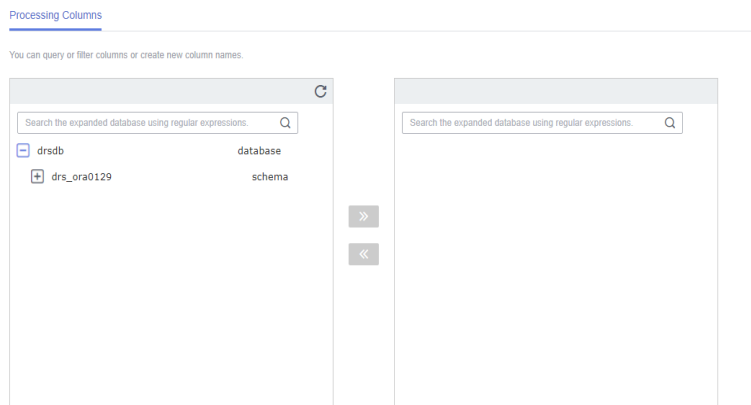
| Parameter | Description |
|------------------------------|---|
| Topic Synchronization Policy | Topic synchronization policy. You can select A specific topic or Auto-generated topics . |
| Topic | Select the topic to be synchronized to the destination database. This parameter is available when the 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. |

| Parameter | Description |
|------------------------|---|
| Synchronize Topic To | <p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database, schema and table names, the performance on a single table query can be improved. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic. This prevents data from being written to the same partition, and consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash value of the database_name.schema.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. |
| Data Format in Kafka | <p>Select the data format to be delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p> |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Figure 4-161 Processing data



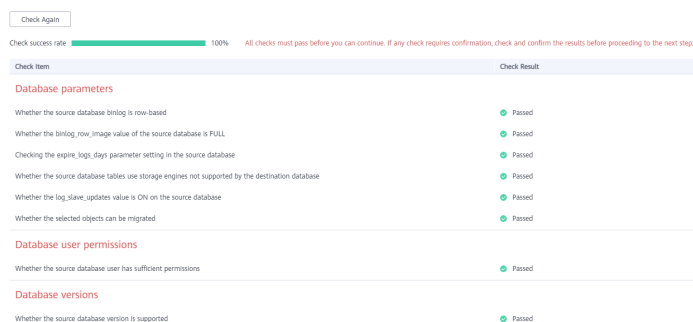
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-162 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-163 Task startup settings





Table 4-174 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.17 From GaussDB Distributed to GaussDB Distributed

Supported Source and Destination Databases

Table 4-175 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 4-176](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-176 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Constraints |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-177](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source 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 4-177 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|--|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|--|------------------|
| | | enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB . | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • To use value comparison, you must have the MONADMIN permission. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-178 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"/?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records |

| Type | Constraints |
|------|---|
| | <p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - 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. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|--|
| 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. |
| 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 GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |

| Type | Constraints |
|----------------------------|--|
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • The topology structure of the source database must be the same as that of the destination database. Otherwise, the content comparison function is unavailable. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-164 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-8131 [Help icon]

Description: [Text area, 0/256 characters]

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 4-179 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-165 Synchronization instance details

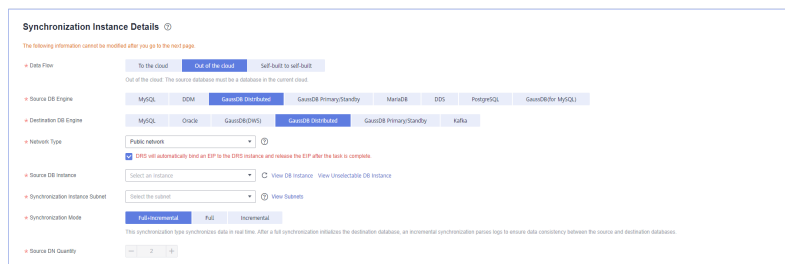


Table 4-180 Synchronization instance settings

| Parameter | Description |
|-----------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Distributed . |
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | <p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The GaussDB distributed instance you created. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPC or VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |
| Source DN Quantity | The value must be the same as the number of DNs in the distributed source database. |

- Task type

Figure 4-166 Task type




Table 4-181 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-167 Enterprise projects and tags




Table 4-182 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-168 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-183 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 4-169 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 4-184 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 4-170 Synchronization mode

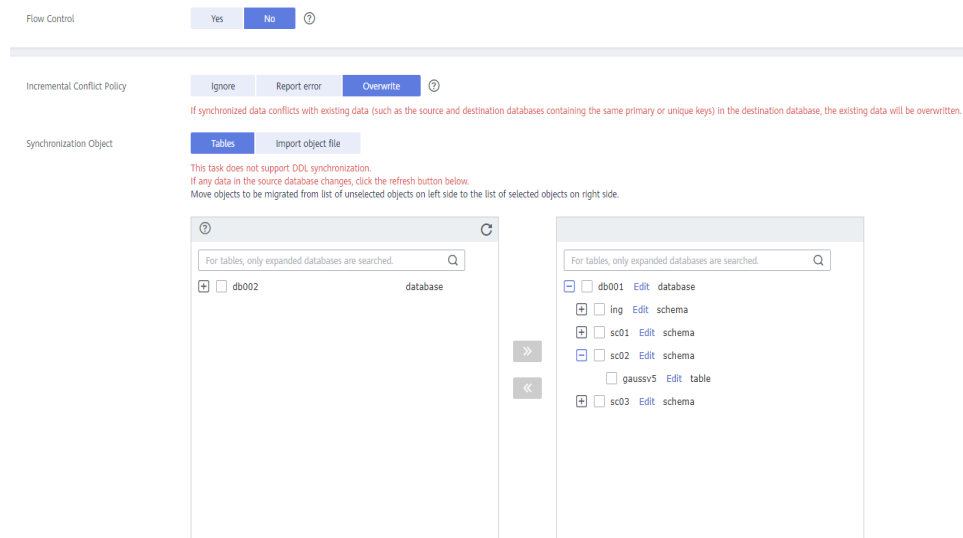
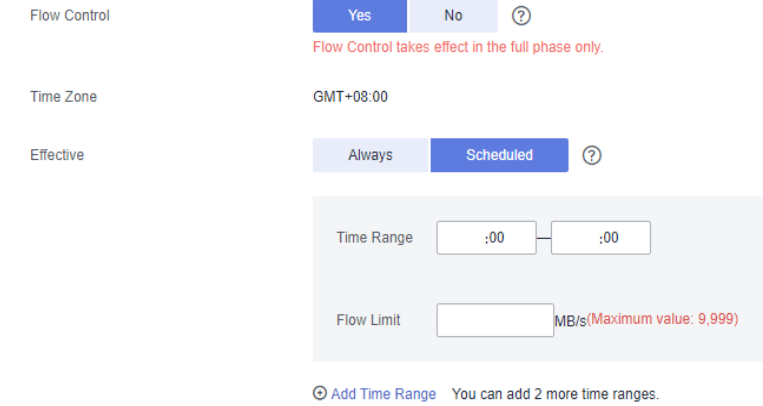



Table 4-185 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-171 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

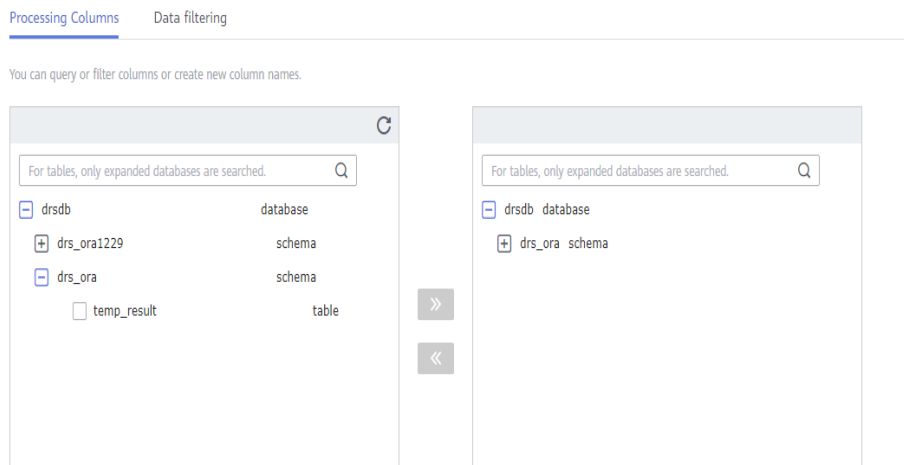
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-172 Processing data



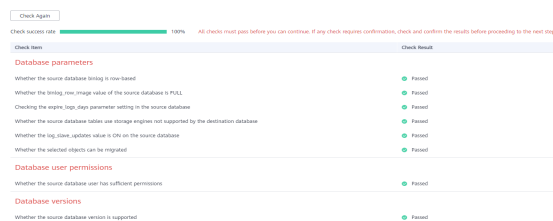
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-173 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-174 Task startup settings





Table 4-186 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.18 From GaussDB Distributed to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 4-187 Supported databases

| Source DB | Destination DB |
|---------------------|---|
| GaussDB distributed | <p>GaussDB primary/standby</p> <p>NOTE The destination database version must be the same as or later than the source database version.</p> |

Supported Synchronization Objects

Table 4-188 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-188 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Constraints |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-189](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- If a distributed database is the source database, the connectivity of each DN needs to be checked. You are advised to perform data synchronization as a non-root database user to prevent user locking due to incorrect password during DN connection.

Table 4-189 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|------------------|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-190 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"/?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records |

| Type | Constraints |
|------|---|
| | <p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - 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. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|---|
| 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. |
| Incremental synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |

| Type | Constraints |
|----------------------------|--|
| Synchronization comparison | <ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-175 Synchronization task information

Table 4-191 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-176 Synchronization instance details

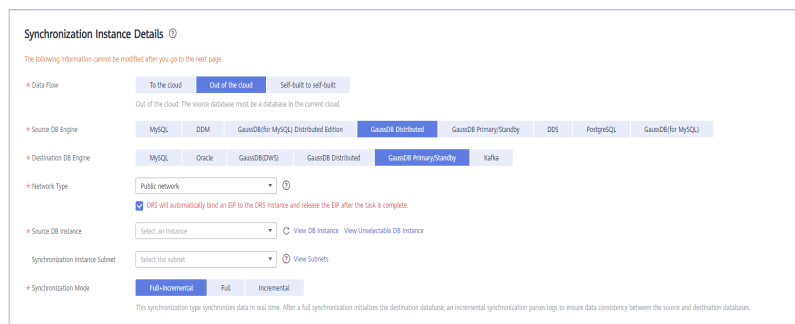


Table 4-192 Synchronization instance settings

| Parameter | Description |
|---------------------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Distributed . |
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | <p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The GaussDB distributed instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-177 Task type



Table 4-193 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-178 Enterprise projects and tags

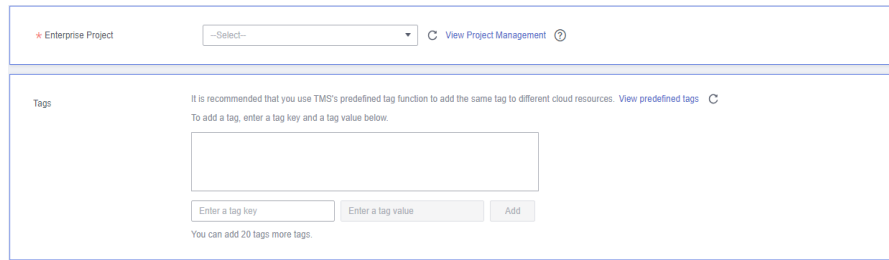


Table 4-194 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-179 Source database information

Table 4-195 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 4-180 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



Test successful

Table 4-196 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 4-181 Synchronization mode

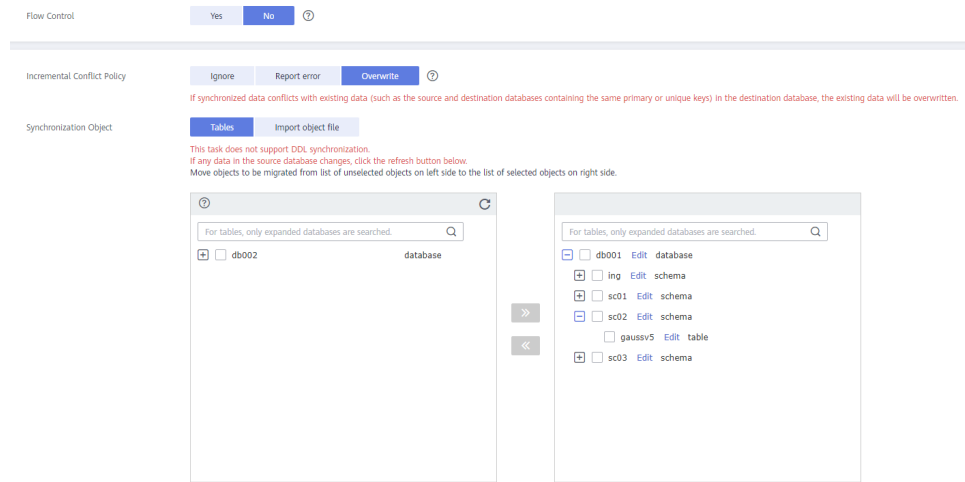
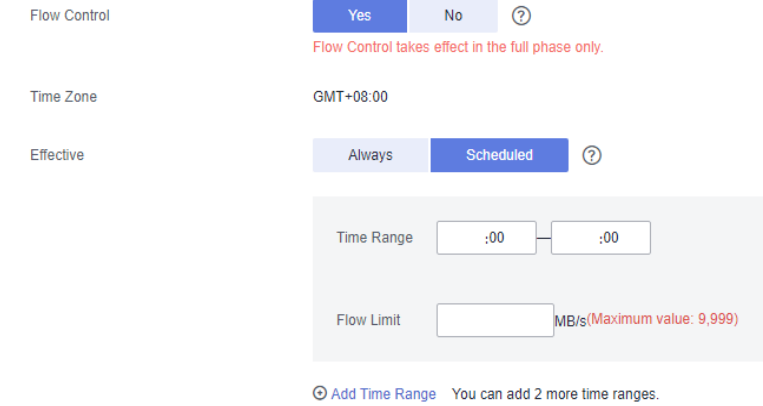



Table 4-197 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-182 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

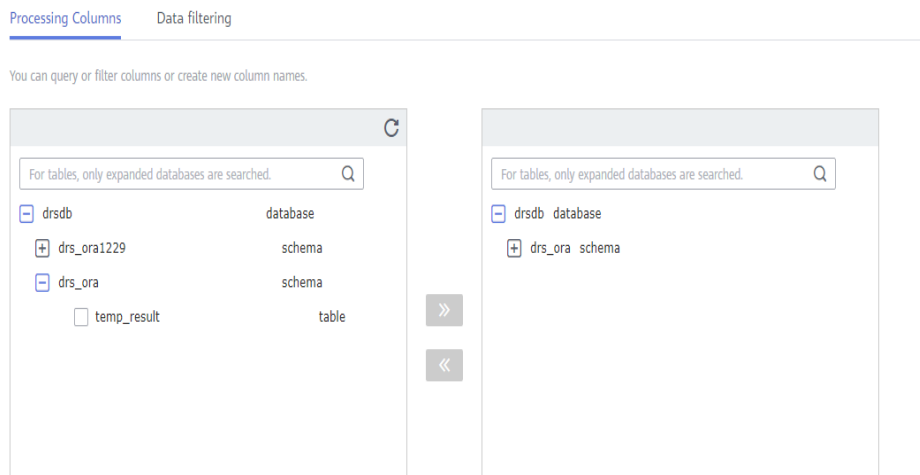
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-183 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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-184 Pre-check

Check success rate 100% All checks must pass before you can continue. If any check requires confirmation, check and confirm the results before proceeding to the next step.

| Check Item | Check Result |
|---|-------------------------|
| Database parameters | |
| Whether the source and destination databases are compatible with each other | Confirm Confirm Details |
| Whether the source database table contains unsupported field types | Passed |
| Source table replication attribute check | Passed |
| Whether the source and destination database character sets are consistent | Passed |
| Whether the destination database contains the configured databases | Passed |
| Whether the target database users(schema) and table exist | Passed |
| Whether the selected table contains delay constraints | Passed |
| Whether the source database contains unlogged tables | Passed |

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-185 Task startup settings

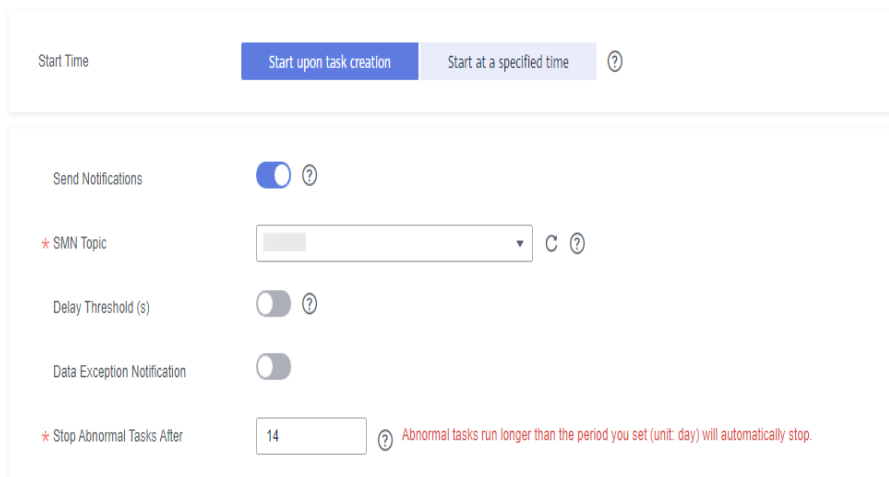



Table 4-198 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.19 From GaussDB Primary/Standby to MySQL

Supported Source and Destination Databases

Table 4-199 Supported databases

| Source DB | Destination DB |
|-------------------------|--|
| GaussDB primary/standby | <ul style="list-style-type: none"> RDS for MySQL 5.6 and 5.7 On-premises MySQL 5.5, 5.6, and 5.7 databases MySQL 5.5, 5.6, and 5.7 databases on an ECS MySQL 5.5, 5.6, and 5.7 databases on other clouds |

Supported Synchronization Objects

Table 4-200 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-200 Supported synchronization objects

| Type | Synchronization Scope |
|-----------------------|--|
| Synchronization scope | <ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. • Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, 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. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbse_perf, pkg_service, pkg_util, dbse_file, dbse_random, dbse_output, dbse_raw, dbse_sql, dbse_lob, dbse_task, blockchain, db4ai, dbse_pldebugger, sqladvisor, dbse_application_info, dbse_match, dbse_pldeveloper, dbse_scheduler, dbse_session, dbse_utility, dbse_sql_util, dbse_xml, dbse_xmldom, dbse_xmlparser, dbse_compression, dbse_heat_map, dbse_ilm, dbse_ilm_admin, prvt_ilm, dbse_profiler, dbse_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 4-201](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-201 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|---|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | <p>a GaussDB Database.</p> <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-202 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\` \?! ● 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 names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - In GaussDB, if there is 0x00 in character or binary data, the data will be truncated, which is inconsistent with MySQL. As a result, the final data may be inconsistent. - Binary data cannot be used as a DML matching condition. That is, binary data cannot be used as a primary key. Binary data is fixed in length and is automatically filled in. The fill character of GaussDB is 0x20, and that of MySQL is 0x00. |
| 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 GaussDB Developer Guide. |
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Do not limit the synchronization speed during data comparison. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. |

| Type | Restrictions |
|-----------------|---|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-186 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

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 4-203 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-187 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: **MySQL** | 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.

* Source 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 4-204 Synchronization instance settings

| Parameter | Description |
|-----------|----------------------------------|
| Data Flow | Select Out of the cloud . |

| Parameter | Description |
|---------------------------------|---|
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select MySQL . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-188 Task type



Table 4-205 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-189 Enterprise projects and tags




Table 4-206 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-190 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-207 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 4-191 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 4-208 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-192 Synchronization mode

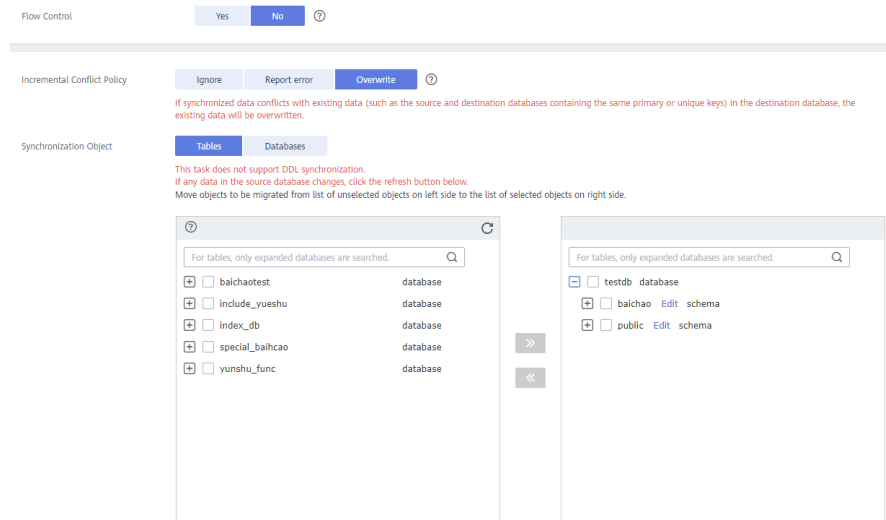
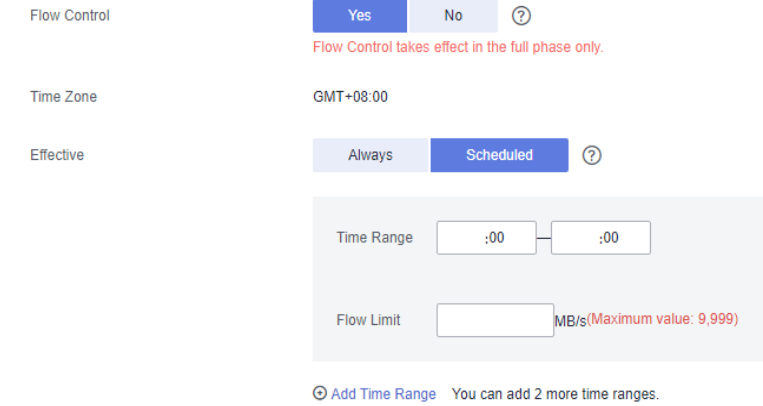



Table 4-209 Synchronization Object

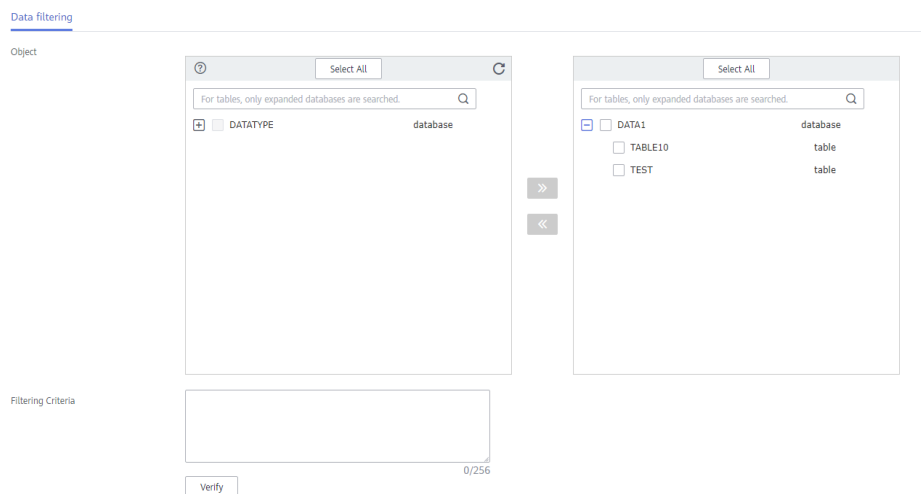
| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-193 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-194 Processing data



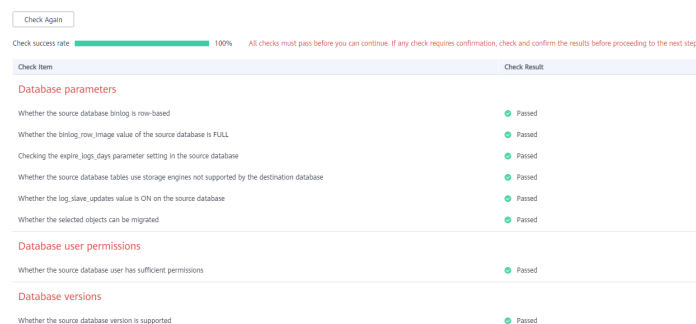
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-195 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-196 Task startup settings





Table 4-210 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.20 From GaussDB Primary/Standby to Oracle

Supported Source and Destination Databases

Table 4-211 Supported databases

| Source DB | Destination DB |
|-------------------------|--|
| GaussDB primary/standby | <ul style="list-style-type: none"> On-premises Oracle 11g and 19c databases Oracle 11g and 19c databases on an ECS |

Supported Synchronization Objects

[Table 4-212](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-212 Supported synchronization objects

| Type | Synchronization Scope |
|-----------------------|---|
| Synchronization scope | <ul style="list-style-type: none"> • Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. • Supported 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 and some DDL statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - 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 4-213](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-213 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|---|--|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role. | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 4-214 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. - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails. - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task |

| Type | Restrictions |
|----------------------|--|
| | <p>failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs.</p> <ul style="list-style-type: none"> - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 change the REPLICIA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding of GaussDB Developer Guide". ● 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. ● The following types of DDL statements cannot be synchronized: ALTER TABLE MERGE PARTITIONS and CREATE INDEX CONCURRENTLY. ● To synchronize DDL statements, ensure that the compatibility modes of the source and destination databases are the same. ● For DDL statements involving tablespace operations, the user of the destination database must have the tablespace operation permissions. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

| Type | Restrictions |
|-----------------|---|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-197 Synchronization task information

Table 4-215 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-198 Synchronization instance details

Table 4-216 Synchronization instance settings

| Parameter | Description |
|-----------|----------------------------------|
| Data Flow | Select Out of the cloud . |

| Parameter | Description |
|---------------------------------|---|
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select Oracle . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full In this mode, data is synchronized from the source to the destination at a time. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-199 Task type



Table 4-217 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-200 Enterprise projects and tags




Table 4-218 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).


- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-201 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

✔ Test successful

Table 4-219 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 4-202 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 4-220 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> • The maximum size of a single certificate file that can be uploaded is 500 KB. • If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-203 Synchronization mode

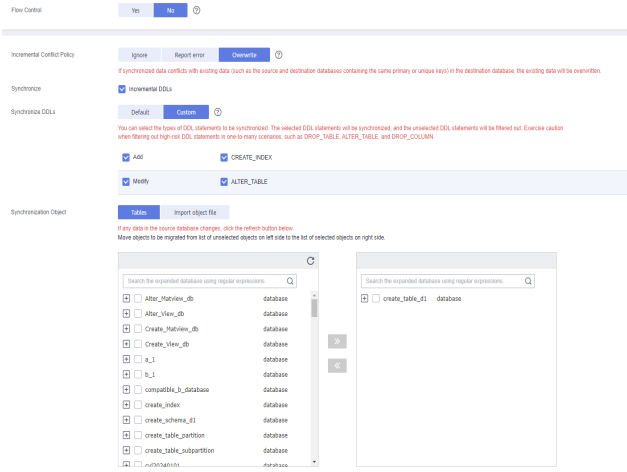
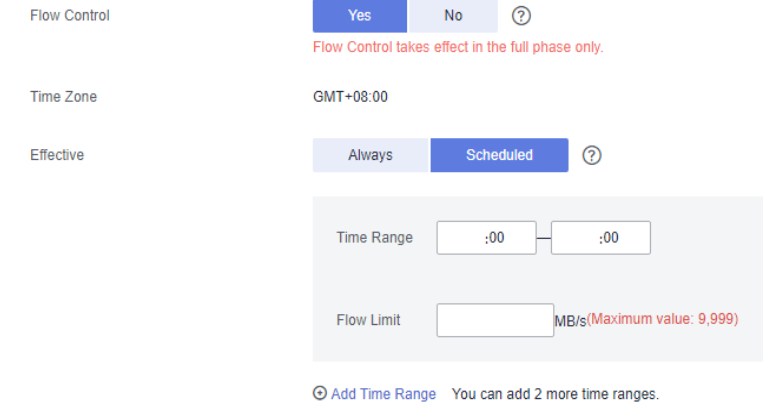



Table 4-221 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-204 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

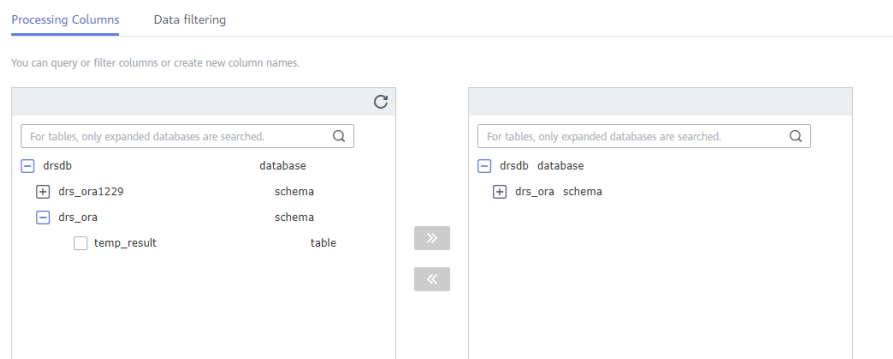
| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| Synchronize | <p>Indicates whether to synchronize DDLs in the incremental synchronization phase.</p> |
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default: DDL statements of the selected synchronization objects will be synchronized to the destination database. <ul style="list-style-type: none"> – DDL statements supported by table-level synchronization: ALTER TABLE and CREATE INDEX ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. Only selected DDL types can be synchronized. |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

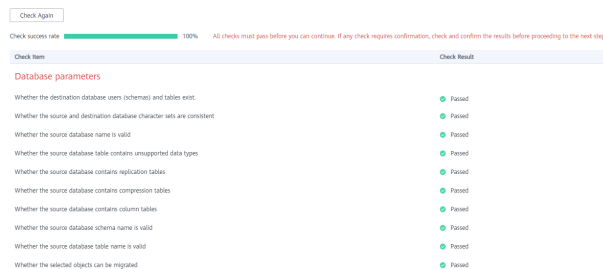
Figure 4-205 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 4-206 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-207 Task startup settings

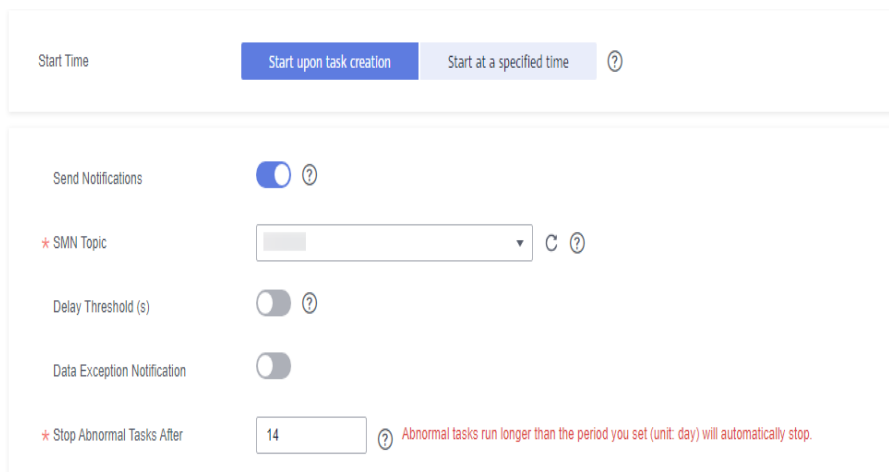



Table 4-222 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |

| Parameter | Description |
|-----------------------------|--|
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.21 From GaussDB Primary/Standby to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-223 Supported databases

| Source DB | Destination DB |
|-------------------------|----------------------|
| GaussDB primary/standby | GaussDB(DWS) cluster |

Supported Synchronization Objects

Table 4-224 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-224 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, DRS allows you to synchronize DML statements of selected tables only. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, and level-2 partitioned tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsibt and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name |

| Type | Constraints |
|------|---|
| | <p>cannot contain double quotation marks ("), single quotation marks ('), or periods (.).</p> <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-225](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-225 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, and SELECT or UPDATE (required only for tables without primary keys) permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|--|--|
| | | <p>Connection and Port Description for Incremental Synchronization from GaussDB.</p> | <p>a GaussDB Database.</p> <ul style="list-style-type: none"> • Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | <p>The user must have the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions on databases • The USAGE and CREATE permissions on schemas • The INSERT, DELETE, UPDATE, and SELECT permissions on tables | <p>The user must have the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> • The CONNECT permission on databases • The USAGE permission on schemas • The INSERT, DELETE, UPDATE, and SELECT permissions on tables | <p>The user must have the sysadmin role or the following permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions on databases • The USAGE and CREATE permissions on schemas • The INSERT, DELETE, UPDATE, and SELECT permissions on tables |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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"> - When a schema name or table name is mapped, to prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key. The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: c_+hash value+original constraint name (which may be truncated)+_key. - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

| Type | Constraints |
|-----------------|---|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-208 Synchronization task information

Table 4-226 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-209 Synchronization instance details

Table 4-227 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 | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |
| Source DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> Full In this mode, data is synchronized from the source to the destination at a time. Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task type

Figure 4-210 Task type



Table 4-228 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-211 Enterprise projects and tags

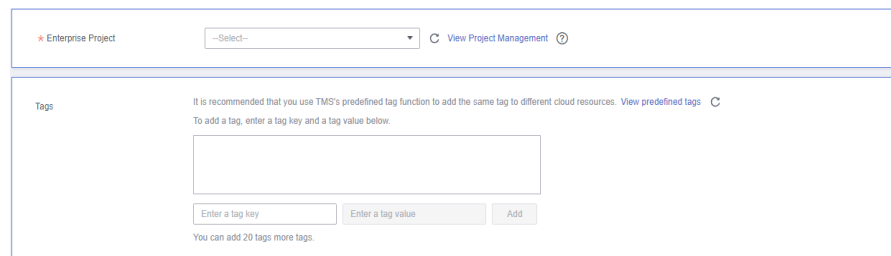


Table 4-229 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-212 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-230 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 4-213 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 4-231 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 4-214 Synchronization mode

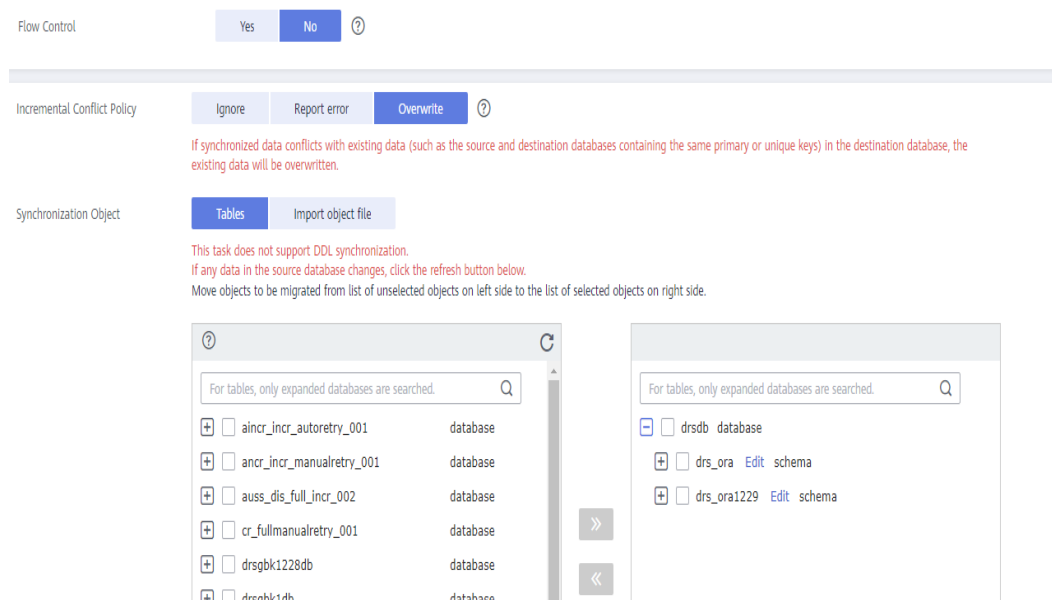
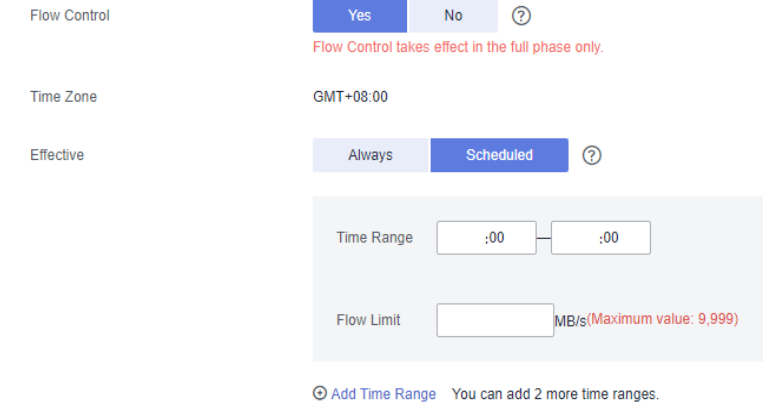



Table 4-232 Synchronization Object


| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-215 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-216 Processing data



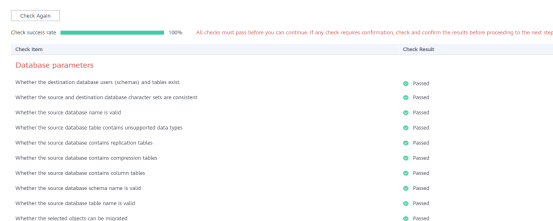
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-217 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 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-218 Task startup settings





Table 4-233 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.22 From GaussDB Primary/Standby to Kafka

Supported Source and Destination Databases

Table 4-234 Supported databases

| Source DB | Destination DB |
|-------------------------|---------------------|
| GaussDB primary/standby | Kafka 0.11 or later |

Supported Synchronization Objects

[Table 4-235](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-235 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 • The supported fields are BIGINT, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE and TINYINT. • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 4-236](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-236 Database user permission

| Type | Incremental |
|----------------------|---|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. • The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. • Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 4-237 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?! ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - If the DRS instance type is set to primary/standby, duplicate or inconsistent data may occur in tables without primary keys when a primary/standby switchover is performed. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|---|
| Incremental synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● During migration of table-level objects, you are not advised to rename the tables. ● Replication of interval partition tables is not supported. ● The name of a primary key column cannot be changed. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> ● After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> ● To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. ● The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-219 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-8131 [Text input]

Description: [Text area]

0/256

Table 4-238 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-220 Synchronization instance details

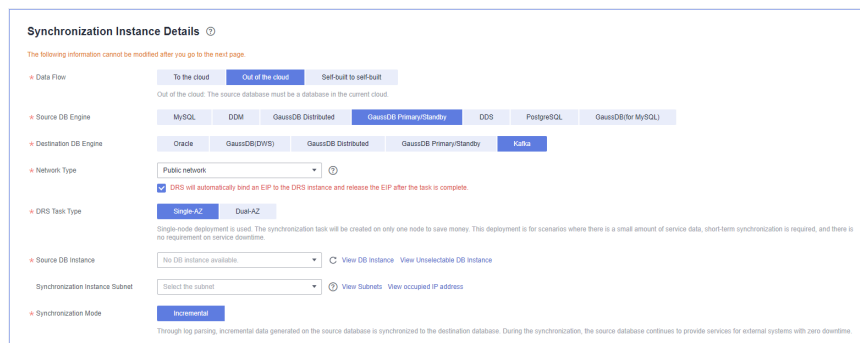


Table 4-239 Synchronization instance settings

| Parameter | Description |
|-----------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select Kafka . |
| Network Type | <p>Public network is used as an example. Available options: Public network and VPN or Direct Connect</p> <ul style="list-style-type: none"> – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases or between cloud databases across regions. |

| Parameter | Description |
|---------------------------------|---|
| DRS Task Type | <p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p> |
| Source DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Synchronization mode | <ul style="list-style-type: none"> - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p> |

- Task Type

Figure 4-221 Task type



Table 4-240 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-222 AZ</p> |

- Enterprise Project and Tags

Figure 4-223 Enterprise projects and tags




Table 4-241 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-224 Source database information

Source Database

DB Instance Name ads-ces-source

Database Username

Database Password

✔ Test successful

Table 4-242 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 4-225 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 4-243 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 4-226 Synchronization mode

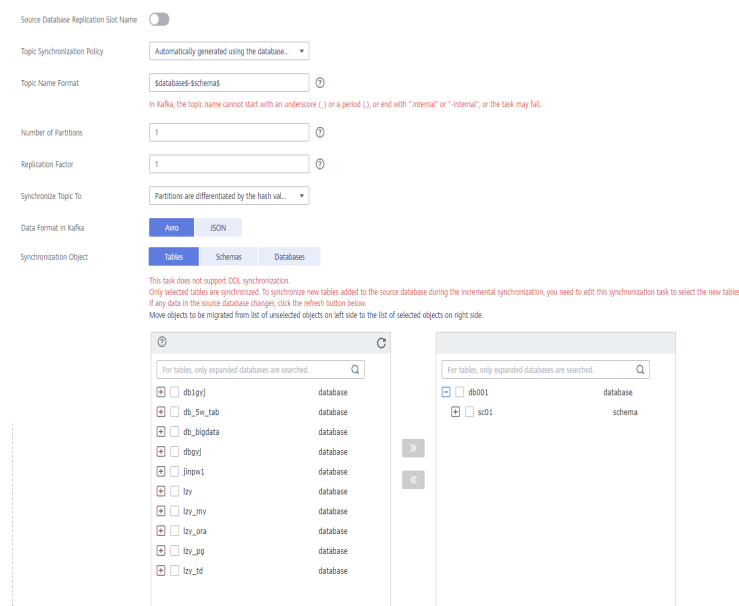



Table 4-244 Synchronization object

| 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. |

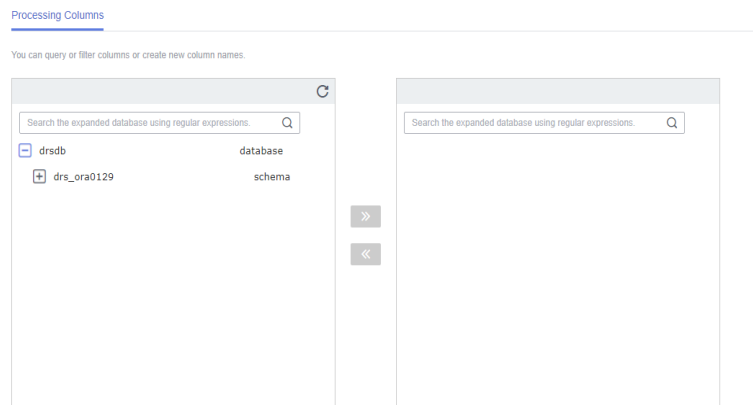
| Parameter | Description |
|------------------------------|---|
| Topic Synchronization Policy | Topic synchronization policy. You can select A specific topic or Auto-generated topics . |
| Topic | Select the topic to be synchronized to the destination database. This parameter is available when Topic Synchronization Policy is set to A specified topic . |
| Topic Name Format | This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal. If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements. |
| Number of Partitions | This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers. |
| Replication Factor | This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created. |

| Parameter | Description |
|------------------------|---|
| Synchronize Topic To | <p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • Partitions are differentiated by the hash values of database_name.schema_name.table_name: This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved. • Partitions are identified by the hash values of the primary key: This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of database_name.schema_name.table_name. • Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. |
| Data Format in Kafka | <p>Select the format of data delivered to Kafka.</p> <ul style="list-style-type: none"> • Avro: A binary encoded format that is efficient. • JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p> |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Figure 4-227 Processing data



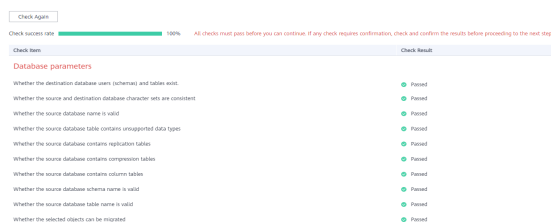
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-228 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-229 Task startup settings





Table 4-245 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.23 From GaussDB Primary/Standby to GaussDB Distributed

Supported Source and Destination Databases

Table 4-246 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 4-247](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-247 Supported synchronization objects

| Type | Restrictions |
|-----------------------|---|
| 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The source table cannot contain a UNIQUE constraint. Otherwise, an error will be reported when data is written to the table structure of GaussDB Distributed, resulting in a task failure. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none">- The database name, schema name, and table name cannot contain special characters /<.>\\' ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.).- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-248](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-248 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|------------------|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-249 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Constraints |
|----------------------|--|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● 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. - If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |
| 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 | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-230 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-6131 [Help icon]

Description: [Text area] 0/256 [Help icon]

Table 4-250 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-231 Synchronization instance details

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 | EDM | 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.

* 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: Full-Incremental | Full | Incremental
This synchronization type synchronizes data in real time. After a full synchronization initializes the destination database, an incremental synchronization passes logs to ensure data consistency between the source and destination databases.

Table 4-251 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | <p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none">- VPC is suitable for data synchronization between cloud databases of the same account in the same region.- Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPC or VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |

| 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. |

- Task type

Figure 4-232 Task type



Table 4-252 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-233 Enterprise projects and tags




Table 4-253 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-234 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-254 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 4-235 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 4-255 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 4-236 Synchronization mode

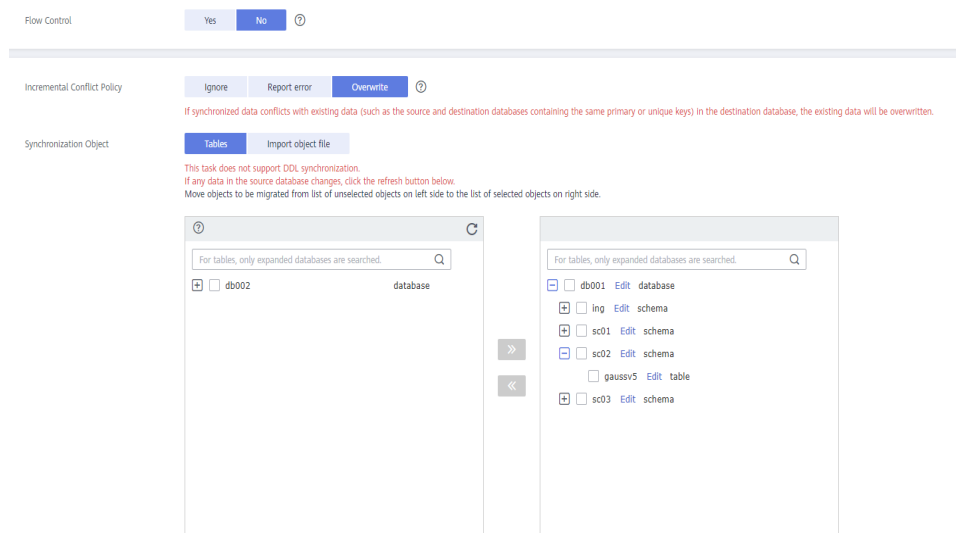
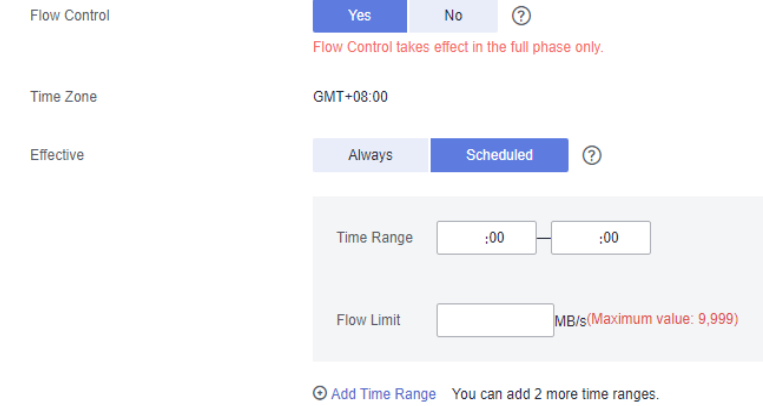



Table 4-256 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-237 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

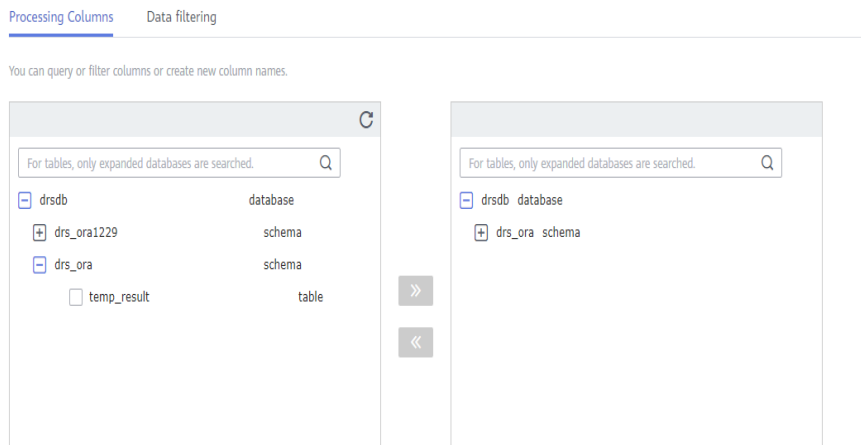
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-238 Processing data



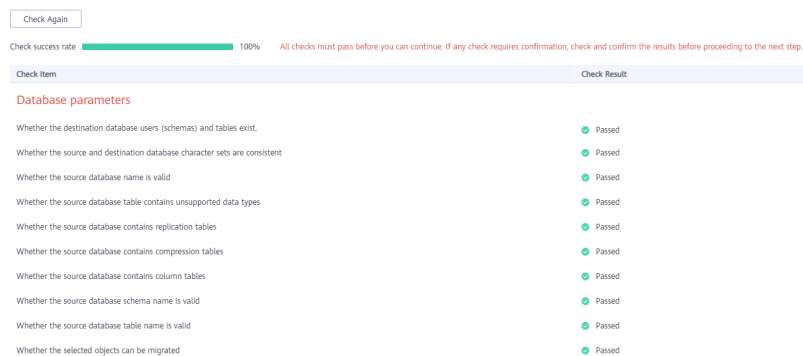
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-239 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-240 Task startup settings





Table 4-257 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.24 From GaussDB Primary/Standby to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 4-258 Supported databases

| Source DB | Destination DB |
|-------------------------|---|
| GaussDB primary/standby | <p>GaussDB primary/standby</p> <p>NOTE The destination database version must be the same as or later than the source database version.</p> |

Supported Synchronization Objects

Table 4-259 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-259 Supported synchronization objects

| Type | Restrictions |
|------------------------------|--|
| Synchronizati on scope | <ul style="list-style-type: none"> ● Instance-level synchronization is not supported. Only one database can be synchronized at a time. To synchronize multiple databases, create multiple tasks. ● Supported scenarios: full synchronization, incremental synchronization, and full+incremental synchronization. ● Supported fields: INTEGER, TINYINT, SMALLINT, BIGINT, NUMBER, NUMERIC, REAL, DOUBLE PRECISION, CHARACTER, CHARACTER VARYING, NVARCHAR2, BIT, BIT VARYING, BLOB, BYTEA, CLOB, RAW, TEXT, JSON, BOOLEAN, DATE, SMALLDATETIME, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE, INTERVAL, BOX, CIDR, CIRCLE, INET, LSEG, MACADDR, MONEY, PATH, POINT, POLYGON, TSQUERY, TSVECTOR, REFCURSOR, UUID and ARRAY ● Table-level synchronization or object file import is supported. <ul style="list-style-type: none"> - During full synchronization, only the structure, sequence, data, constraints, and indexes of the selected table can be synchronized. - During incremental synchronization, only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be selected. - Column-store tables, compressed tables, delay tables, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System catalogs (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-260](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-260 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|------------------|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 4-261 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Select full or full+incremental synchronization. If a table structure already exists in the destination database, ensure that the table structure is the same as that on the source database. If column processing is performed, ensure that the table structure is the same as that after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Constraints |
|----------------------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● 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. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |
| 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 | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 4-241 Synchronization task information

Table 4-262 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-242 Synchronization instance details

Table 4-263 Synchronization instance settings

| Parameter | Description |
|-----------|----------------------------------|
| Data Flow | Select Out of the cloud . |

| Parameter | Description |
|---------------------------------|---|
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | <p>Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The GaussDB primary/standby instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPC or VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |

| 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. |

- Task type

Figure 4-243 Task type



Table 4-264 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-244 Enterprise projects and tags




Table 4-265 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-245 Source database information

Source Database

DB Instance Name rds-cn-sou

Database Username

Database Password

✔ Test successful

Table 4-266 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 4-246 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 4-267 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 4-247 Synchronization mode

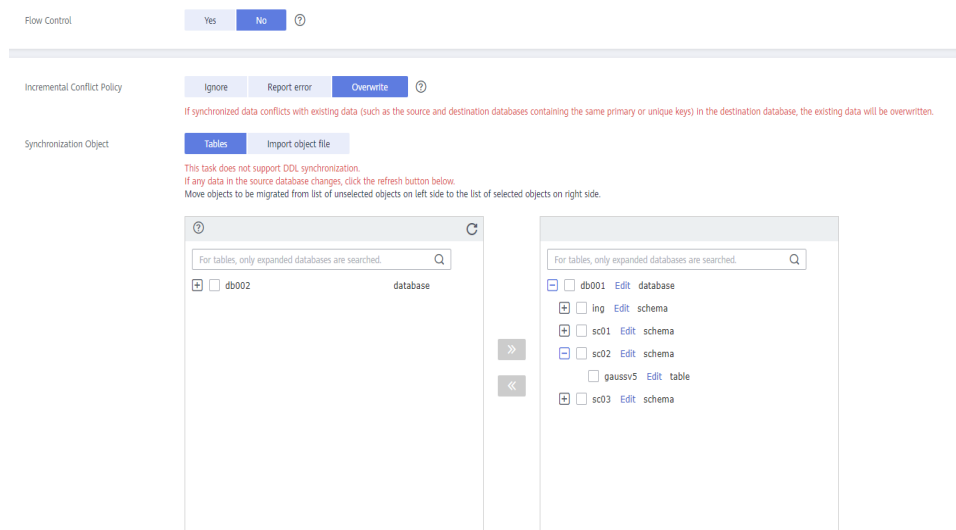
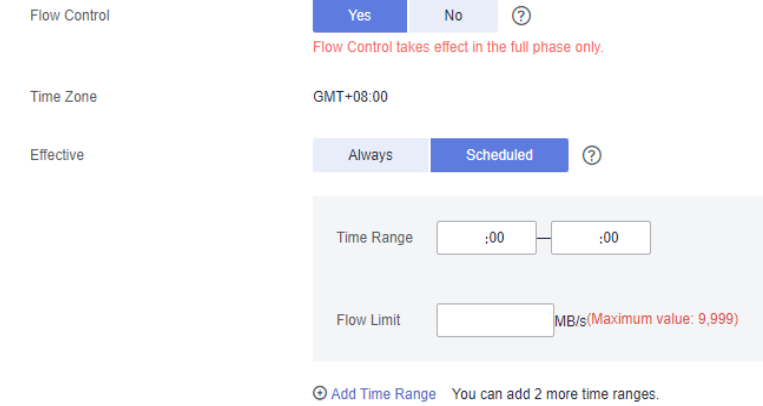



Table 4-268 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-248 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

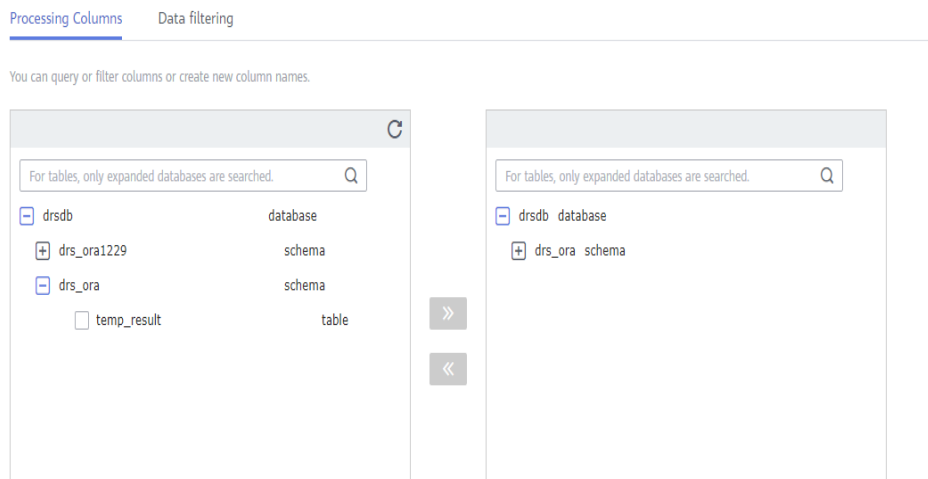
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 4-249 Processing data



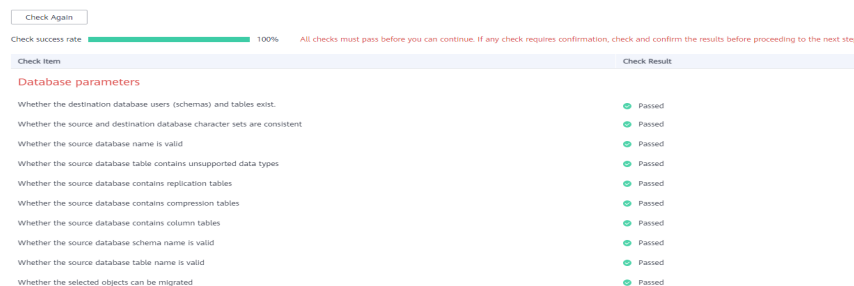
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-250 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-251 Task startup settings





Table 4-269 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.25 From GaussDB(for MySQL) to MySQL

Supported Source and Destination Databases

Table 4-270 Supported databases

| Source DB | Destination DB |
|---|--|
| <ul style="list-style-type: none"> GaussDB(for MySQL) Primary/ Standby | <ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-271 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. - Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|---|
| Source database | <ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the MySQL source database server_id must be set. The value ranges from 1 to 4294967296. ● The names of the source databases, tables, and views cannot contain non-ASCII characters, or the following characters: '<>/' ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized to RDS. ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● Database mapping does not support views or stored procedures. If the source database contains views or stored procedures, the synchronization may fail. |
| Destination database | <ul style="list-style-type: none"> ● The destination database is running properly. ● The destination database must have sufficient disk space. ● If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. ● The character set of the destination database must be the same as that of the source database. ● The time zone of the destination database must be the same as that of the source database. ● If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source or destination database, the DRS task may fail. ● If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss. ● The source and destination databases cannot contain tables that have the same names but do not have primary keys. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● The source database does not support point-in-time recovery (PITR). ● The destination database cannot be restored to a point in time when a full synchronization was being performed. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization. ● DDL operations are not supported during full synchronization. ● During incremental synchronization, some DDL operations are supported. <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, CHANGE_COLUMN, DROP_COLUMN, DROP_INDEX, ADD_INDEX, CREATE_INDEX, RENAME_INDEX, DROP_TABLE, TRUNCATE_TABLE, DROP_PARTITION, RENAME_COLUMN, DROP_PRIMARY_KEY and ADD_PRIMARY_KEY. You can select the DDL operations to be synchronized on the object selection page as required. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> - Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • If you use additional columns and the number of columns in a single table exceeds 500, adding additional columns may lead to the number of columns in a table to reach an upper limit and cause the task to fail. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-252 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, the form contains the following 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 help icon.

 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." The bottom right corner of the form area shows the number "0/256".

Table 4-272 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |

| Parameter | Description |
|-------------|--|
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-253 Synchronization instance details

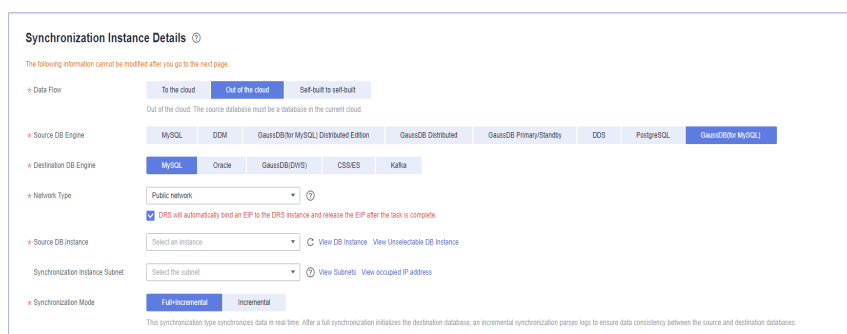


Table 4-273 Synchronization instance settings

| Parameter | Description |
|-----------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB(for MySQL) . |
| Destination DB Engine | Select MySQL . |
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |

| Parameter | Description |
|---------------------------------|---|
| Source DB Instance | The GaussDB(for MySQL) instance you created. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Synchronization Mode | <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. |

- Task type

Figure 4-254 Task type




Table 4-274 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-255 Enterprise projects and tags

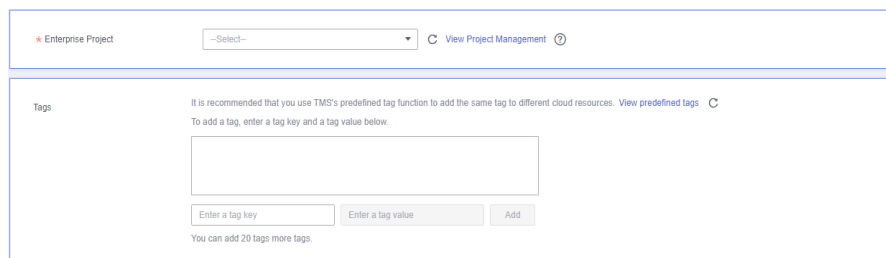


Table 4-275 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-256 Source database information

Source Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 4-276 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 4-257 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 4-277 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 4-258 Synchronization mode

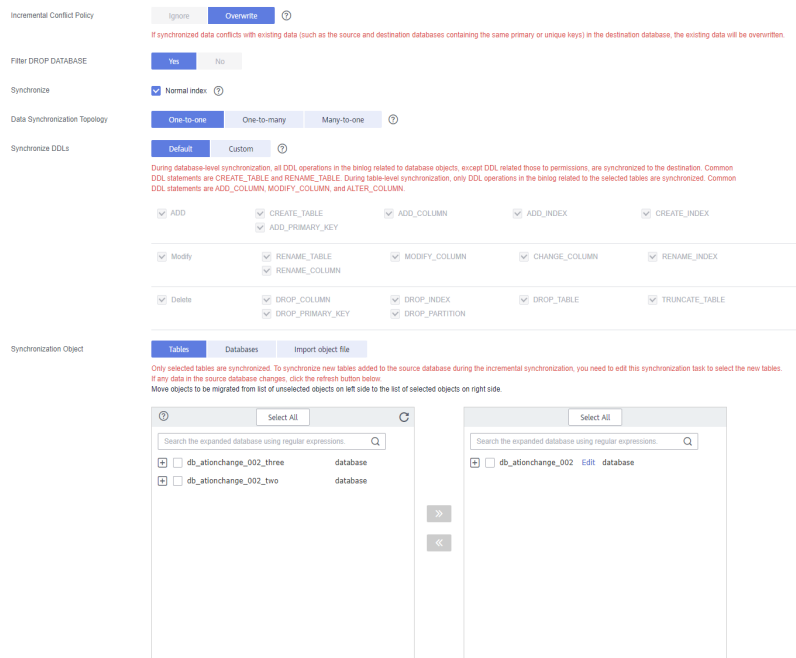



Table 4-278 Synchronization mode and object

| Parameter | Description |
|-------------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">• Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database.• Overwrite Conflicting data will be overwritten. |
| 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. |
| Start Point | <p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| Data Synchronization Topology | <p>Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p> <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>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. |

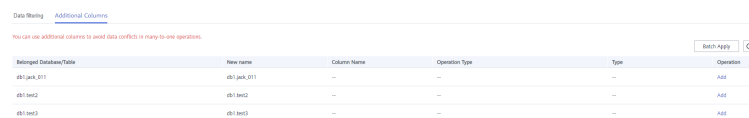
| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: 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. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If you need to set data filtering, click **Data Filtering** and set related filtering rules.
- If you need to add additional columns, click the **Additional Columns** tab, click **Add** in the **Operation** column, and enter the column name and the operation type.

For details about related operations, see [Processing Data](#).

Figure 4-259 Processing data



| Selected Database/Table | New Name | Column Name | Operation Type | Type | Operation |
|-------------------------|--------------|-------------|----------------|------|-----------|
| db1_jack_011 | db1_jack_011 | -- | -- | -- | Add |
| db1_0002 | db1_0002 | -- | -- | -- | Add |
| db1_0003 | db1_0003 | -- | -- | -- | Add |

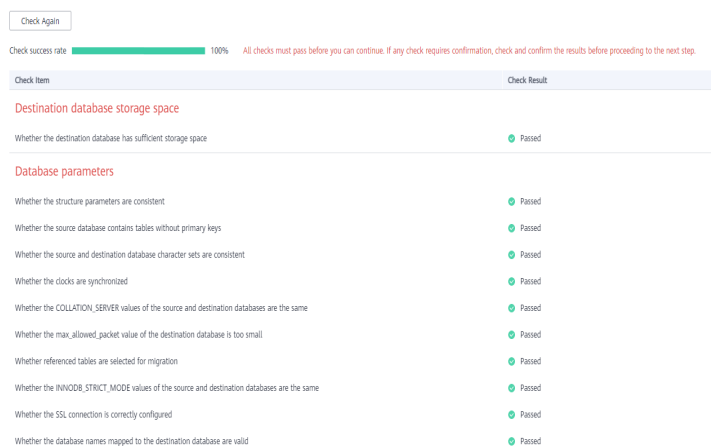
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-260 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-261 Task startup settings

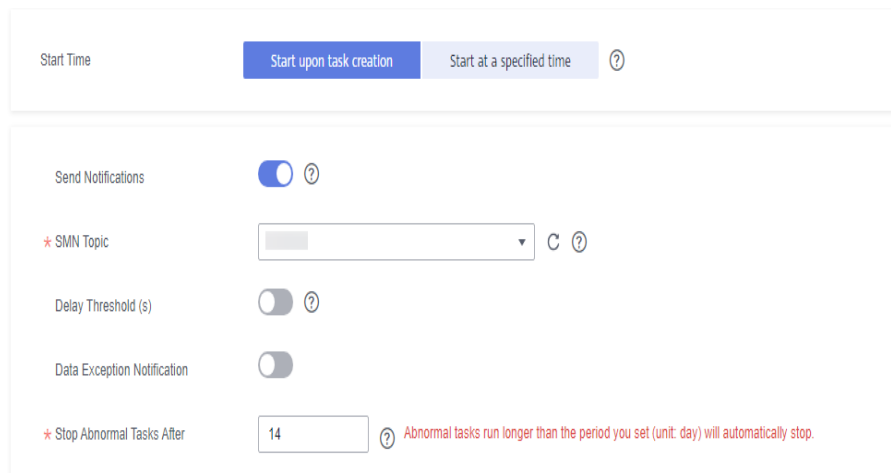



Table 4-279 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.26 From GaussDB(for MySQL) to GaussDB(DWS)

Supported Source and Destination Databases

Table 4-280 Supported databases

| Source DB | Destination DB |
|--|----------------------|
| Primary/standby GaussDB(for MySQL) instances | GaussDB(DWS) cluster |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-281 Precautions

| Type | Restrictions |
|----------------------|--|
| Database permissions | <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. • The destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, and CREATE. |

| Type | Restrictions |
|------------------------|--|
| Synchronization object | <ul style="list-style-type: none"> ● Only tables, common indexes (B-Tree indexes), and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, virtual columns, unique constraints, and unique indexes cannot be synchronized. ● Comment is supported in full synchronization mode, but not supported in incremental synchronization mode. ● The following data types are not supported: XML, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, and multipolygon. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● The column name of the synchronization object cannot be a field forbidden by GaussDB(DWS), such as CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID and TID. Otherwise, the task fails. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |
| Source database | <ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● The 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 '<>/\ |

| Type | Restrictions |
|----------------------|--|
| Destination database | <ul style="list-style-type: none">• The destination database is running properly.• The destination database must have sufficient disk space.• The time zone of the destination database must be the same as that of the source database. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common association: Indexes reference tables. ● If the network is reconnected within 30 seconds, real-time synchronization will not be affected. If the network is interrupted for more than 30 seconds, the synchronization task will fail. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB(DWS), the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format after synchronization: hash value + original constraint name (which may be truncated) + _key, or table name_original index name. ● In the full synchronization phase, only B-Tree indexes are synchronized. Other indexes are not synchronized by default. If there are too many GaussDB(DWS) indexes, the storage space and data import performance will be affected. You are advised to create indexes based on service requirements. ● If there is a unique key when a table without a primary key is synchronized, data may fail to be written due to data conflicts. ● If the character sets of the source and destination databases are different, data may be inconsistent or synchronization may fail. ● If the data types are incompatible, the synchronization may fail. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • During real-time synchronization, you can change the port number. If the synchronization task fails after the port is changed, you can retry the synchronization task. • During real-time synchronization, the IP address, account, and password cannot be changed. • During incremental synchronization, the following DDL operations are supported: <ul style="list-style-type: none"> - In one-to-one synchronization, the following DDL operations are synchronized by default: CREATE_TABLE, RENAME_TABLE, ADD_COLUMN, MODIFY_COLUMN, ALTER_COLUMN, DROP_CONSTRAINT, ADD_CONSTRAINT, CREATE_INDEX, DROP_INDEX, RENAME_INDEX, and RENAME_COLUMN. You can select the DDL operations to be synchronized on the object selection page as required. - If you rename a column in many-to-one synchronization, you must stop services. Otherwise, data inconsistency may occur. - In many-to-one scenarios, you are advised to synchronize ADD_COLUMN only. Other DDL synchronization may cause task failures or data inconsistency due to destination table changes. - In many-to-one scenarios, when running ADD_COLUMN, ensure that the types of columns added to each table are the same. Otherwise, the task may fail. - The name of a table, column, or index to be added or modified cannot exceed 63 characters. Otherwise, the task fails. - When an index is creating using a DDL statement, if the table name and index name in the statement contain more than 63 characters, duplicate names may occur. As a result, the index fails to be created. - If a primary key is added to a table that does not have a primary key in the source database, the DDL operation must contain the first column. Otherwise, the task fails. - When a DDL operation is performed in the incremental phase, if the destination table is not found, the DDL operation will be ignored. - In the incremental phase, if CHANGE COLUMN is performed in the source database to modify a column and the column is a distribution column in destination |

| Type | Restrictions |
|------|--|
| | <p>GaussDB(DWS), the statement will be ignored because GaussDB(DWS) does not support distribution column modification.</p> <ul style="list-style-type: none"> - In the incremental phase, the RENAME INDEX operation is not supported because the index rules of the source and destination databases are different. - In the incremental phase, you are not advised to combine CHAR(0) with other characters for data synchronization. For example, inserting CONCAT('a',CHAR(0),'b') may cause data inconsistencies. - In the incremental phase, you are not advised to combine CHAR(34) (double quotation marks) and CHAR(92) (backslash) with other characters for JSON data synchronization. For example, inserting JSON_OBJECT('\$\$.0', CONCAT('a', CHAR(34), 'b')) may cause data inconsistencies because an extra backslash (\) will be added after it is synchronized to the destination end. <ul style="list-style-type: none"> • During an incremental synchronization, database-level synchronization does not support online DDL, and table-level synchronization supports only online DDL generated by Alibaba Cloud DMS. • During an incremental synchronization, tables whose primary key type is binary, text, blob, or clob cannot be deleted or updated. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-262 Synchronization task information

Table 4-282 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-263 Synchronization instance details

Table 4-283 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 | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | An available GaussDB(for MySQL) instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |

- Task type

Figure 4-264 Task type




Table 4-284 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-265 Enterprise projects and tags

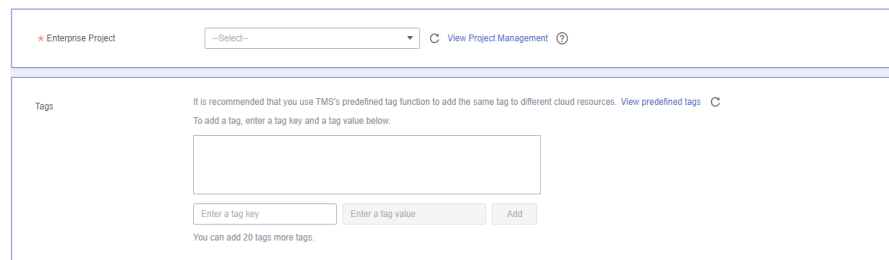


Table 4-285 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-266 Source database information

Source Database

DB Instance Name ods-oms-source

Database Username

Database Password

✔ Test successful

Table 4-286 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 4-267 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 4-287 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 4-268 Synchronization mode

Flow Control Yes No ?

Synchronization Object Type Table structure Data Index
When you manually create a table structure in the destination database, for details about the data type, see Mapping Data Types.

Incremental Conflict Policy Ignore Report error Override ?
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.

Synchronize DML Insert Update Delete ?

Data Synchronization Topology One-to-one One-to-many Many-to-one ?

Synchronize DDLs Default Custom ?
During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN.

| | | | | |
|--|---|--|--|--|
| <input checked="" type="checkbox"/> ADD | <input checked="" type="checkbox"/> CREATE_TABLE | <input checked="" type="checkbox"/> ADD_COLUMN | <input checked="" type="checkbox"/> ADD_CONSTRAINT | <input checked="" type="checkbox"/> CREATE_INDEX |
| <input checked="" type="checkbox"/> Modify | <input checked="" type="checkbox"/> RENAME_TABLE <input checked="" type="checkbox"/> CHANGE_COLUMN | <input checked="" type="checkbox"/> MODIFY_COLUMN <input checked="" type="checkbox"/> RENAME_COLUMN | <input checked="" type="checkbox"/> ALTER_COLUMN | <input checked="" type="checkbox"/> RENAME_INDEX |
| <input checked="" type="checkbox"/> Delete | <input checked="" type="checkbox"/> DROP_CONSTRAINT <input checked="" type="checkbox"/> TRUNCATE_TABLE | <input checked="" type="checkbox"/> DROP_INDEX <input checked="" type="checkbox"/> DROP_PARTITION | <input checked="" type="checkbox"/> DROP_COLUMN | <input checked="" type="checkbox"/> DROP_TABLE |

Synchronization Object Tables Databases Import object file
Only some DDL statements can be synchronized. For details, see precautions of the current scenario in Real-Time Synchronization > Before You Start.
 If any data in the source database changes, click the refresh button below.
 Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

? Select All ↻

Search the expanded database using regular expressions.

- db_doublezucstargetecs_002 database
- db_doublezucstargetecs_003 database
- db_kernel001ssync_kernel_001 database
- db_m2o_001_4216 database
- db_ransfer_vpc_001 database
- db_ransfer_vpc_001_definer database
- db_snnrc_rebuild_002 database
- db_snnrc_rebuild_002_three database
- db_snnrc_rebuild_002_two database
- db_tabase_001 database
- dba_dp1dataprocess_001 database
- dba_dp2dataprocess_001 database

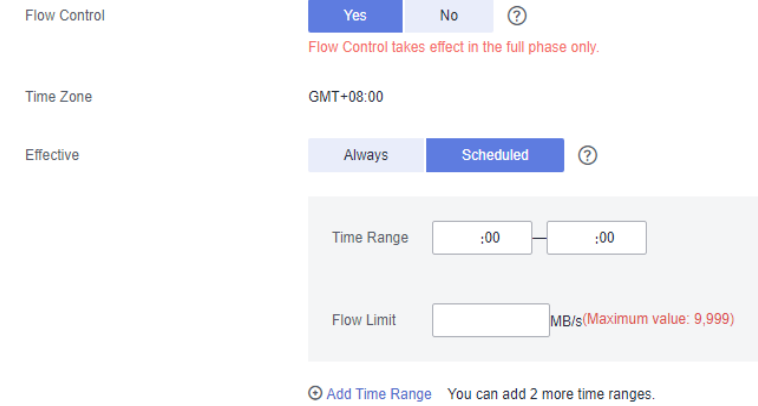
? Select All

Search the expanded database using regular expressions.

database ?


gtest Edit

Table 4-288 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-269 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-------------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> |
| Synchronize DML | <p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> |
| Data Synchronization Topology | <p>Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies.</p> <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>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"> In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |


Step 5 On the **Data Processing** page, select the table object to be processed, click **Add** in the **Operation** column, enter the column name, type, and operation type, confirm the information, and click **Next**. You can configure related rules by referring to [Processing Data](#).

Figure 4-270 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.

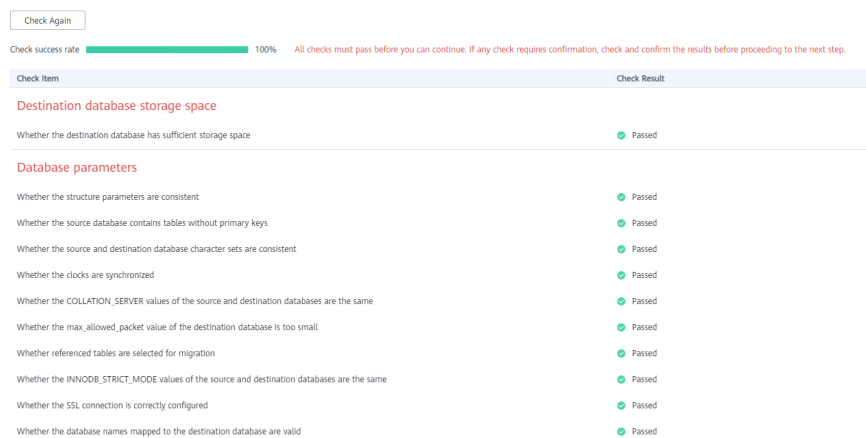
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-271 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-272 Task startup settings

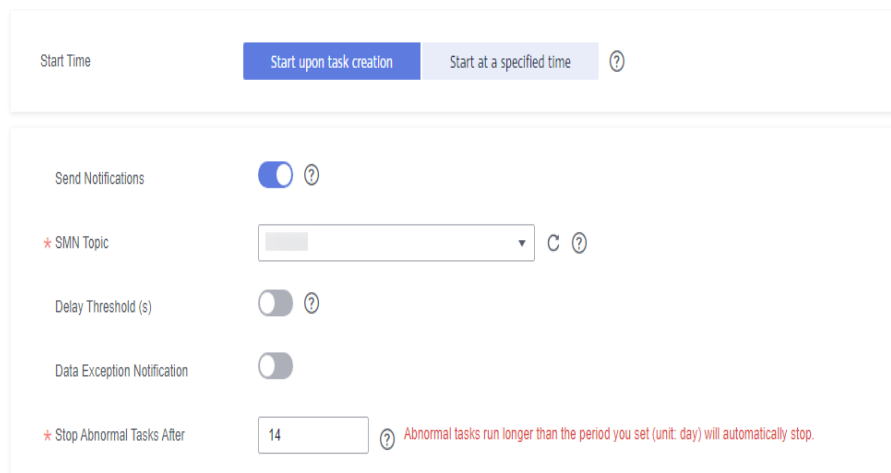



Table 4-289 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none">• If the delay threshold is set to 0, no notifications will be sent to the recipient.• In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent.• Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none">• You can set this parameter only for pay-per-use tasks.• Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.27 From GaussDB(for MySQL) to Kafka

Supported Source and Destination Databases

Table 4-290 Supported databases

| Source DB | Destination DB |
|--|----------------|
| Primary/standby GaussDB(for MySQL) instances | Kafka |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
- To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - Tables to be synchronized without a primary key may be locked for 3s.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-291 Precautions

| Type | Restrictions |
|------------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. |
| Synchronization object | <ul style="list-style-type: none"> • During full synchronization, tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized, but events and triggers cannot be synchronized. During incremental synchronization, only table data and DDLs can be synchronized. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● 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 '<'>/'\ |
| Destination database | <ul style="list-style-type: none"> ● The destination database is a Kafka database. ● You are advised to set auto.create.topics.enable of Kafka to false. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● If the data types are incompatible, the synchronization may fail. ● If a full synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the identifier field in the Kafka data for data deduplication. (The shard ID must be unique.) ● During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database does not support point-in-time recovery (PITR). ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During synchronization of table-level objects, renaming tables is not recommended. ● During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase. |

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 4-273 Synchronization task information

Table 4-292 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-274 Synchronization instance details

Table 4-293 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 Kafka . |
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| DRS Task Type | <p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p> |
| Source DB Instance | The GaussDB(for MySQL) instance you created. |

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>Available options: Full+Incremental and Incremental</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <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. |


- Task Type

Figure 4-275 Task type



Table 4-294 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |

| Parameter | Description |
|-----------|---|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-276 AZ</p>  |

- Enterprise Project and Tags

Figure 4-277 Enterprise projects and tags




Table 4-295 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-278 Source database information

Source Database

DB Instance Name gaussdb-for-mysql-1

Database Username

Database Password

✔ Test successful

Table 4-296 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 4-279 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

Table 4-297 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 4-280 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.

C

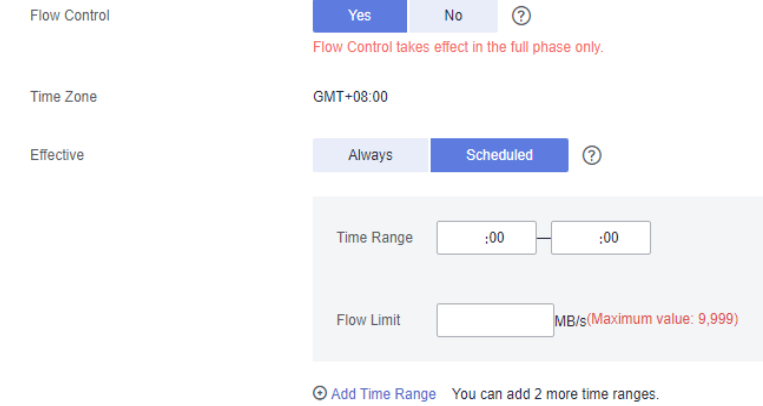
For tables, only expanded databases are searched. Q

- baichao database
- db_taurus_omsprecheck_001 database
- dengyuhuan database
- dengyuhuan1 database
- drs database
- drs4 database
- dyh1 database
- dyh1209 database


C

For tables, only expanded databases are searched. Q

Table 4-298 Synchronization Object

| Parameter | Description |
|-----------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-281 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 | Select the DML operations to be synchronized. By default, all DML operations are selected. |

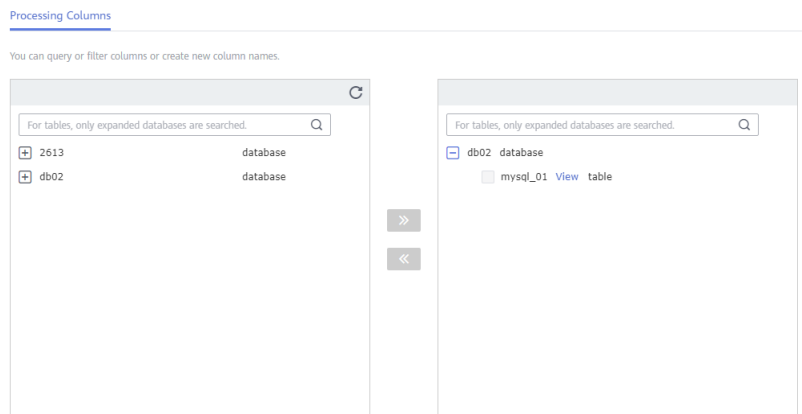
| Parameter | Description |
|------------------------------|---|
| Start Point | <p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| Topic Synchronization Policy | <p>Topic synchronization policy. You can select A specific topic or Auto-generated topics.</p> |
| Topic | <p>Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic.</p> |
| 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 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. |

| Parameter | Description |
|------------------------|--|
| Data Format in Kafka | <p>Select the data format to be delivered from GaussDB(for MySQL) to Kafka.</p> <ul style="list-style-type: none"> ● Avro refers to binary encoded format. This option is available only when Synchronization Mode is set to Incremental in Step 2. Only whitelisted users can use the Avro option. To use this option, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket. ● JSON: JSON message format, which is easy to interpret but takes up more space. ● JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 4-282 Processing data



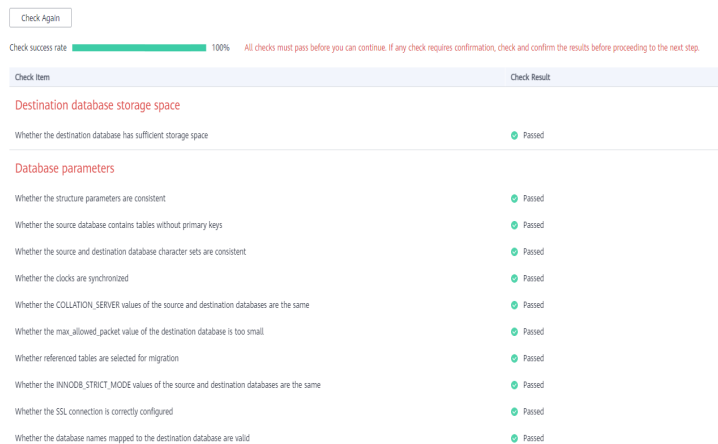
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-283 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-284 Task startup settings





Table 4-299 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.28 From GaussDB(for MySQL) to CSS/ES

Supported Source and Destination Databases

Table 4-300 Supported databases

| Source DB | Destination DB |
|--|---|
| <ul style="list-style-type: none"> Primary/standby GaussDB(for MySQL) instances | <ul style="list-style-type: none"> ElasticSearch 5.5, 6.2, 6.5, 7.1, 7.6, 7.9 and 7.10 |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-301 Precautions

| Type | Restrictions |
|------------------------|--|
| Database permissions | <p>Minimum permission requirements for full plus incremental synchronization:</p> <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT • The destination database user must have the following permissions: READ and WRITE |
| Synchronization object | <ul style="list-style-type: none"> • The table data can be synchronized. • Databases, views, indexes, constraints, functions, stored procedures, triggers, and events cannot be synchronized. • The system database and event status cannot be synchronized. • Tables whose primary keys are of the FLOAT type cannot be synchronized. • Tables that do not have primary keys cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|---|
| Source database | <ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● The source database names cannot contain non-ASCII characters, or the following characters: '<>/' ● The table name in the source database cannot contain non-ASCII characters or the following characters: '<>/' ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● During an incremental synchronization, the server_id value of the source database must be specified. The value of server_id ranges from 1 to 4294967296. ● Enable skip-name-resolve for the source database to reduce the possibility of connection timeout. ● Enable the Global Transaction Identifier (GTID) of the source database. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails. ● The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail. ● The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail. |
| Destination database | <ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● The source database cannot be restored to a point in time when a full synchronization was being performed. ● Convert the value range of the source database time field that is not supported by the destination database to null. ● The strings padded with zeros in the source database may be truncated because the source database uses the fixed-length binary data type, and the destination database uses the variable-length data type. ● The binary value is encrypted using Base64 and then written to the destination database. ● If no time zone is specified for the source database, specify the time zone for synchronizing the datetime type to the destination database. ● All table field names are converted to lowercase letters. ● If the <code>_id</code> field of the destination database is generated using multiple columns in the source database, separate these columns with colons (:). ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. ● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot write data using the statement-based binlog format. ● During the synchronization, do not clear binlogs on the source database. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● During the synchronization, do not create a database named ib_logfile in the source. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> • During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. • Incremental synchronization filters out all DDL operations. • During incremental synchronization, resumable upload is supported, but data may be repeatedly inserted into non-transactional tables that do not have primary keys when the server system breaks down. • If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-285 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

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

* Task Name ⓘ

Description ⓘ

0/256

Table 4-302 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-286 Synchronization instance details

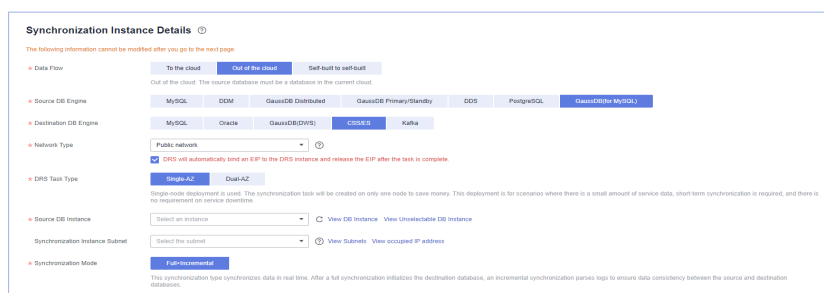


Table 4-303 Synchronization instance settings

| Parameter | Description |
|-----------------------|------------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB(for MySQL) . |
| Destination DB Engine | Select CSS/ES . |

| Parameter | Description |
|--------------------|---|
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| DRS Task Type | <p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p> |
| Source DB Instance | The GaussDB(for MySQL) instance you created. |

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <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. |

- Task Type

Figure 4-287 Task type



Table 4-304 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 4-288 AZ</p> |

- Enterprise Project and Tags

Figure 4-289 Enterprise projects and tags




Table 4-305 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-290 Source database information

Source Database

DB Instance Name rds-csm-source

Database Username

Database Password

✔ Test successful

Table 4-306 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 4-291 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 4-307 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. Only .cer and .pem certificates are supported. If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-292 Synchronization mode

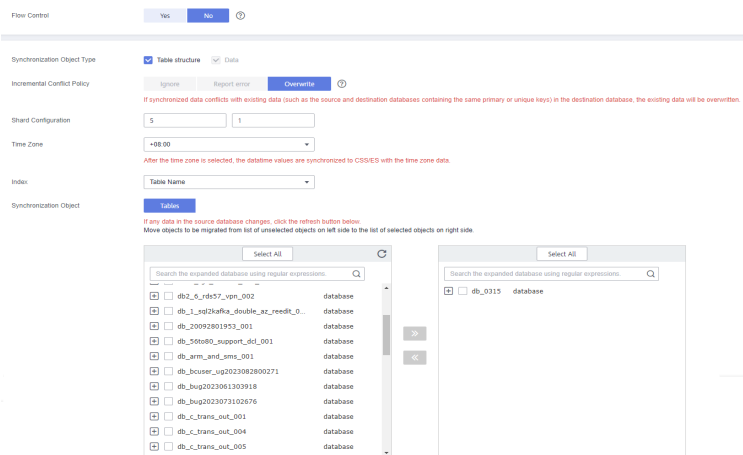
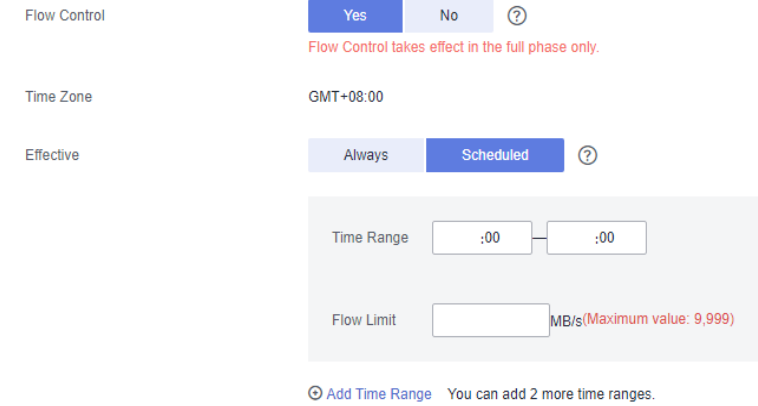



Table 4-308 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-293 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy only applies to incremental synchronization. The default value is Overwrite. The conflict in the full synchronization phase is ignored by default.</p> |
| Shard Configuration | <p>Configure the number of primary shards and the number of replicas. The default number of primary shards is 5, and the default number of shard replicas is 1.</p> |
| Time Zone | <p>After a time zone is selected, the datetime values are synchronized to CSS/ES with the time zone data.</p> |
| Index Name | <ul style="list-style-type: none"> • Table Name The index name created in the target Elasticsearch instance is the same as the table name. • Database Name+Table Name. The name of the index created in the target Elasticsearch instance is <i>DatabaseName_TableName</i>. |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-level synchronization. You can select data for synchronization based on your service requirements.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects. For details, see Editing a Synchronization Task. <p>NOTE</p> <ul style="list-style-type: none"> - To quickly select the desired database objects, you can use the search function. - If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. - If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |


Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering** or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

 **NOTE**

The column processing of the synchronization task from GaussDB(for MySQL) to CSS/ES supports only column filtering and does not support column mapping.

Figure 4-294 Processing data



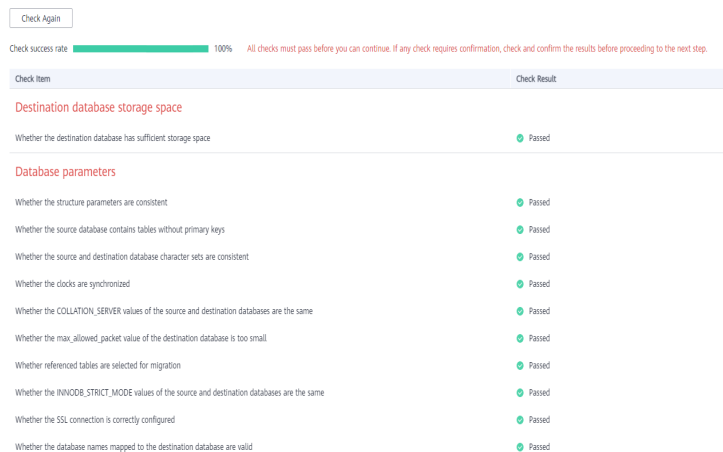
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-295 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-296 Task startup settings





Table 4-309 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.29 From GaussDB(for MySQL) to Oracle

Supported Source and Destination Databases

Table 4-310 Supported databases

| Source DB | Destination DB |
|--|--|
| <ul style="list-style-type: none"> Primary/standby GaussDB(for MySQL) instances | <ul style="list-style-type: none"> On-premises Oracle databases Oracle databases on an ECS |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-311 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. • The maximum number of columns supported by the source GaussDB(for MySQL) database is 1017, while the maximum number of columns supported by the destination Oracle database is 1000. Therefore, the number of columns of objects to be synchronized cannot exceed 1000. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● The source database must be the primary node of the primary/standby GaussDB(for MySQL) instance. ● The source database names cannot contain non-ASCII characters, or the following characters: '<>/' ● The table name in the source database cannot contain non-ASCII characters or the following characters: '<>/' ● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. ● During the incremental synchronization, the binlog of the source database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. ● If the expire_logs_days value of the source database is set to 0, the synchronization may fail. ● During an incremental synchronization, the source database server_id must be set to a value ranging from 1 to 4294967296. ● Enable skip-name-resolve for the source database to reduce the possibility of connection timeout. ● Enable the Global Transaction Identifier (GTID) of the source database. ● The source database does not support the mysql binlog dump command. ● The character set of the source database must be the same as that of the destination database. Otherwise, the synchronization fails. ● The log_slave_updates parameter of the source database must be enabled. Otherwise, the synchronization will fail. ● The binlog_row_image parameter of the source database must be set to FULL. Otherwise, the synchronization will fail. ● Do not set lower_case_table_names to 0. |
| Destination database | <ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space. |

| Type | Restrictions |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● If the default value of the time field in the source database is all 0s, it will be converted to 1970-01-01 00:00:00. ● If the precision of source database decimal data type exceeds 38, the data will be truncated because the maximum precision of the destination database number data type is 38. ● If the length of the varchar data in the source database is greater than or equal to 667 characters, the varchar type will be converted to the clob type in the Oracle database. ● 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. |

| 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. |

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 4-297 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 4-312 Task information

| Parameter | Description |
|-----------|---|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |

| Parameter | Description |
|-------------|--|
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-298 Synchronization instance details

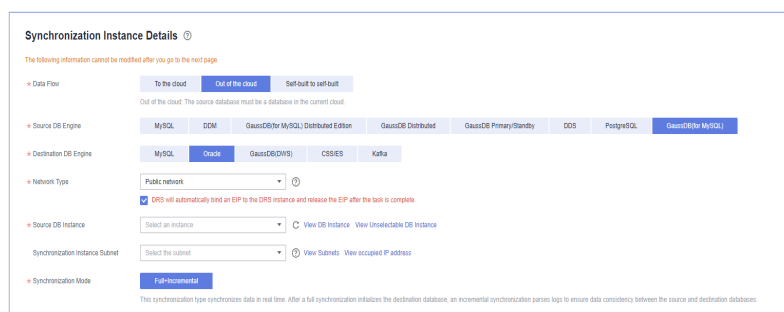


Table 4-313 Synchronization instance settings

| Parameter | Description |
|-----------------------|--|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select GaussDB(for MySQL) . |
| Destination DB Engine | Select Oracle . |
| Network Type | <p>The Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect</p> <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The GaussDB(for MySQL) instance you created. |

| Parameter | Description |
|---------------------------------|--|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <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. |

- Task type

Figure 4-299 Task type



Table 4-314 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> |

- Enterprise Project and Tags

Figure 4-300 Enterprise projects and tags




Table 4-315 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 4-301 Source database information

Source Database

DB Instance Name ods-oms-source

Database Username

Database Password

✔ Test successful

Table 4-316 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 4-302 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 4-317 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address to improve access performance. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-303 Synchronization mode

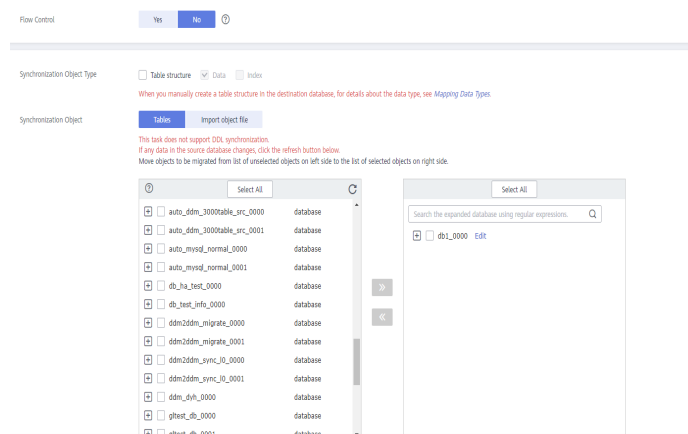
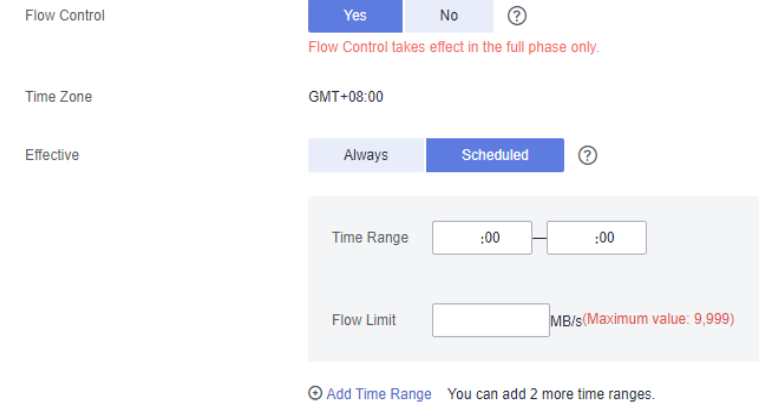



Table 4-318 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-304 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <p>For details about how to import an object file, see Importing Synchronization Objects.</p> <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

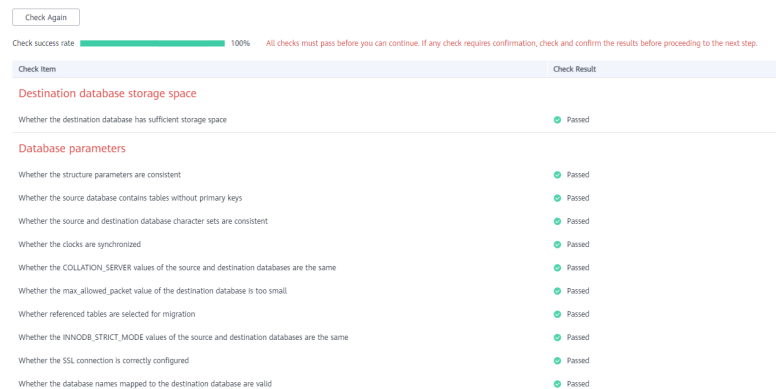
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-305 Pre-check



 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-306 Task startup settings

Table 4-319 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.30 From MariaDB to MariaDB

Supported Source and Destination Databases

Table 4-320 Supported databases

| Source DB | Destination DB |
|--|---|
| <ul style="list-style-type: none">RDS for MariaDB 10.5 | <ul style="list-style-type: none">On-premises MariaDB 10.5ECS-hosted MariaDB 10.5Other cloud MariaDB 10.5 |

Supported Synchronization Objects

Table 4-321 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-321 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> ● Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● Events and triggers cannot be synchronized. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● 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. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent. ● If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur. |

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization

tasks require different permissions. For details, see [Table 4-322](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-322 Database account permission

| Type | Full+Incremental |
|---------------------------|--|
| Source database user | SELECT, SHOW VIEW, EVENT, LOCK TABLES, RELOAD, REPLICATION SLAVE, REPLICATION CLIENT (or BINLOG MONIOTOR, changed to this permission in MariaDB 10.5 and later versions) |
| Destination database user | SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

- If you read a table, especially a large table, during the full 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 4-323 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - During the incremental synchronization, the binlog of the source MariaDB database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. - If the expire_logs_days value of the source database is set to 0, the synchronization may fail. - Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. - During an incremental synchronization, the source MariaDB database server_id must be set to a value ranging from 1 to 4294967296. ● Source database object requirements: <ul style="list-style-type: none"> - The source database names cannot contain non-ASCII characters, or the following characters: '<>/' - 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. ● Destination database parameter requirements: <ul style="list-style-type: none"> - Data cannot be synchronized from a newer version database to an older version database. - The character set of the destination database must be the same as that of the source database. - The time zone of the destination database must be the same as that of the source database. - During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. - If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails. ● Destination database object requirements: |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> - The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal. - The destination DB instance must have sufficient storage space. - If the destination database (excluding the system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database. - The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (') <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. - 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 DB instances are RDS for MariaDB instances, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. - If the source MariaDB database does not support TLS1.2, you need to submit an application to the O&M personnel before using SSL to test the connection. - 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 |

| Type | Restrictions |
|----------------------|--|
| | <p>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"> - The source database does not support point-in-time recovery (PITR). - The destination database cannot be restored to a point in time when a full synchronization was being performed. - Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. - Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. - The partitioned table does not support column mapping. - Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - After a task is created, the destination database cannot be set to read-only. - If index synchronization is required for a DRS task, the destination database table cannot contain indexes with the same name but different columns. In the full phase, DRS ignores the existing indexes with the same name. In the incremental phase, DDL operations on indexes based on the index name trigger misplacement. |
| Full synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Data inconsistency may occur when the MyISAM table is modified during a full synchronization. |

| 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 write data to the destination database. Otherwise, data may be inconsistent. • DDL statements can be synchronized during incremental synchronization. • Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected. • You can add additional objects during an incremental synchronization. |
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison. |

Procedure

This section describes how to use DRS to configure a MariaDB out-of-cloud synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 4-307 Synchronization task information

⚠ Only the task name and description can be modified. Other settings cannot be modified after you click Create Now on this page.
The system will create virtual resources immediately after you click Create Now. Virtual resources cannot be modified after being created so no settings except the task name and description can be modified.

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-6131 [Help icon]

Description: [Text area] 0/256 [Help icon]

Table 4-324 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 4-308 Synchronization instance information

Synchronization Instance Details ⓘ

The following information cannot be modified after you go to the next page.

- Data Flow: To the cloud | **Out of the cloud** | Self built to self built
- Source DB Engine: MySQL | DDM | GaussDB(for MySQL) Distributed Edition | GaussDB Distributed | GaussDB Primary/Standby | **MariaDB** | DDS | PostgreSQL | GaussDB(for MySQL)
- Destination DB Engine: MySQL
- Network Type: Public network
- Source DB Instance: No DB instance available | View DB Instance | View Unavailable DB Instance
- Synchronization Instance Subnet: Select the subnet | View Subnets
- 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.

Table 4-325 Synchronization instance settings

| Parameter | Description |
|------------------|---|
| Data Flow | Select Out of the cloud . The source database is a database on the current cloud. |
| Source DB Engine | Select MariaDB . |

| Parameter | Description |
|---------------------------------|---|
| Destination DB Engine | Select MariaDB . |
| Network Type | <p>Available options: VPC, Public network and VPN or Direct Connect. Public network is the default value and is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | The RDS for MariaDB instance you created. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | <p>Full+Incremental: This synchronization mode allows you to synchronize data without interrupting services. After a full synchronization initializes the destination database, an incremental synchronization initiates and parses logs to ensure data consistency between the source and destination databases.</p> <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> |

- Task type

Figure 4-309 Task type




Table 4-326 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-310 Enterprise projects and tags




Table 4-327 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 4-311 Source database information

The screenshot shows a configuration form titled "Source Database". It contains three input fields: "DB Instance Name" with the value "rds-mariadb-src", "Database Username" with the value "root", and "Database Password" which is masked with dots. Below the fields is a "Test Connection" button, and to its right is a green checkmark icon followed by the text "Test successful".

Table 4-328 Source database information

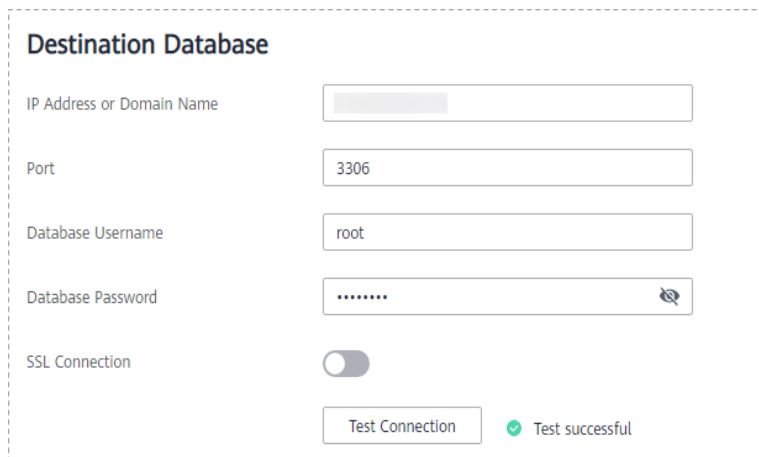
| Parameter | Description |
|-------------------|--|
| DB Instance Name | The RDS for MariaDB instance you selected when creating the task. The parameter cannot be changed. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in the system until the task is deleted.

- Destination database configuration

Figure 4-312 Destination database information




Destination Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 4-329 Source database information

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 4-313 Synchronization objects

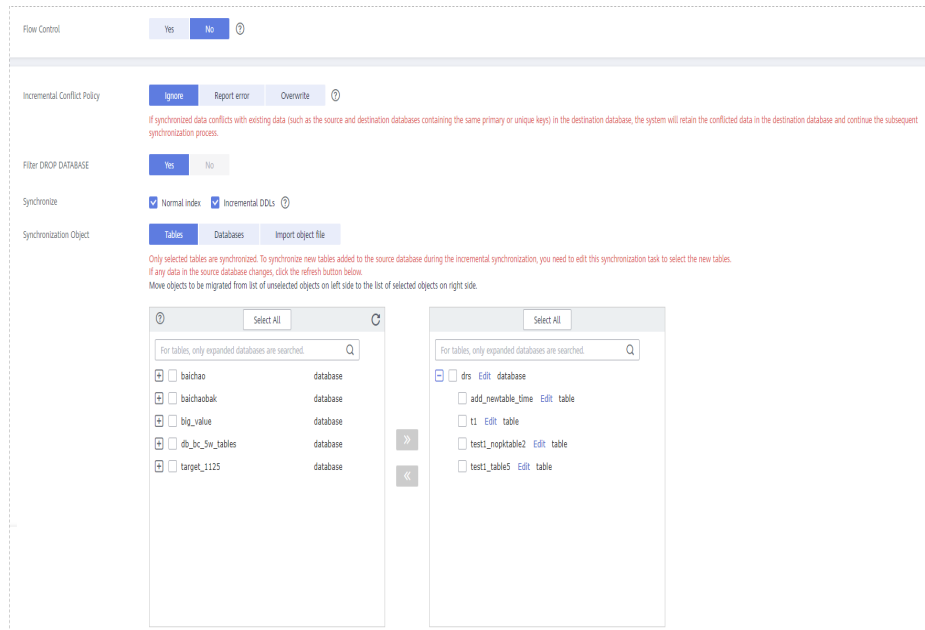
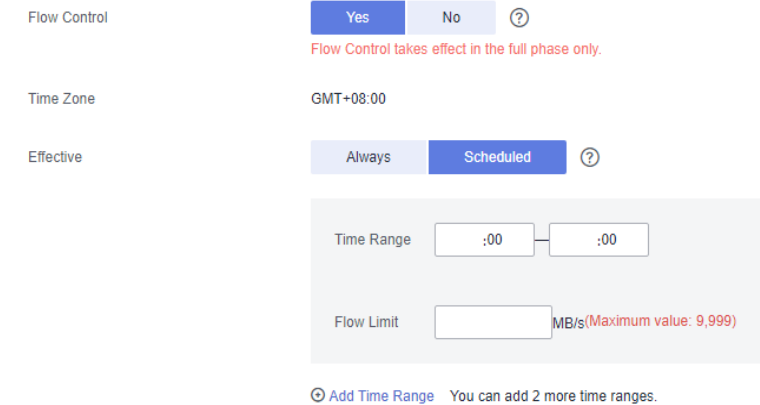



Table 4-330 Synchronization mode and object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 4-314 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> • If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. • If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize data based on service requirements.</p> |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. <ul style="list-style-type: none"> – If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. – In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. • For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**. For details about how to configure related rules, see [Processing Data](#).

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

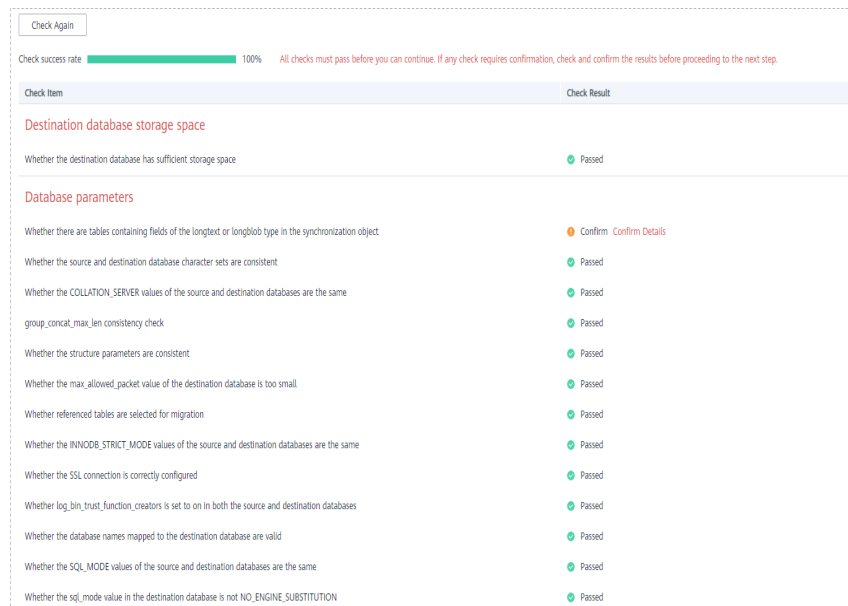
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Figure 4-315 Pre-check



Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 4-316 Task startup settings

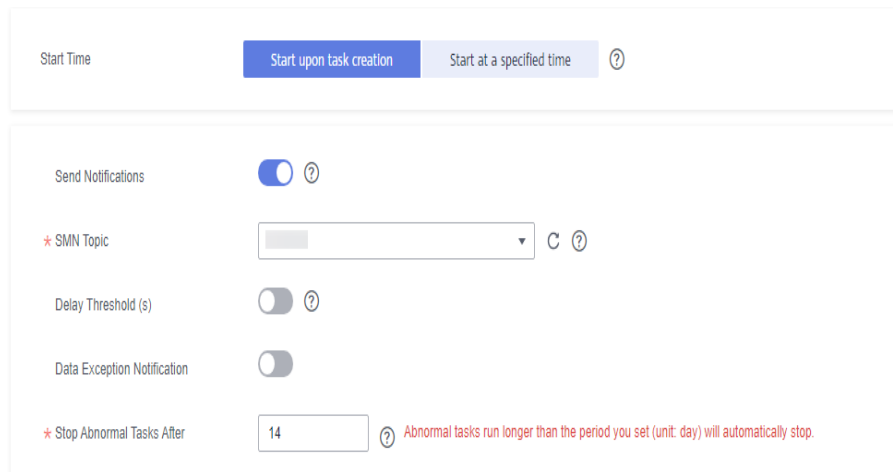



Table 4-331 Task startup settings

| Parameter | Description |
|------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |

| Parameter | Description |
|-----------------------------|--|
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

4.31 From Microsoft SQL Server to Kafka

Supported Source and Destination Databases

Table 4-332 Supported databases

| Source DB | Destination DB |
|--|---------------------|
| RDS for SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) | Kafka 0.11 or later |

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 4-333 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-333 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> • Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER • Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY • Scope of incremental synchronization <ul style="list-style-type: none"> - DML statements, including INSERT, UPDATE, and DELETE, are supported. - DDL statements are not supported. |

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization

tasks require different permissions. For details, see [Table 4-334](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 4-334 Database account permission

| Type | Incremental |
|----------------------|--|
| Source database user | Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized |

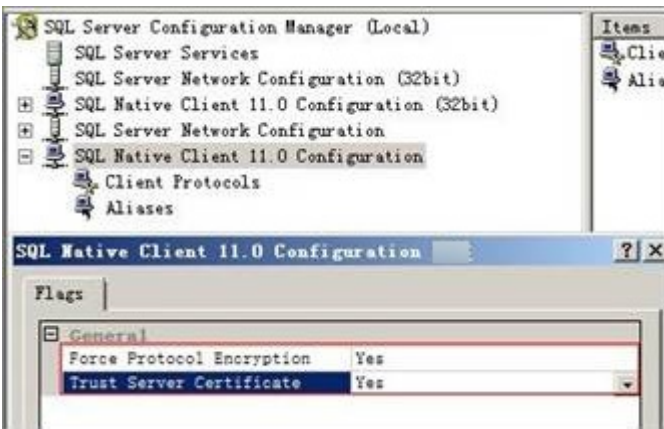
Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 4-335 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - The SQL Server Agent proxy service must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 4-317. <p style="text-align: center;">Figure 4-317 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database parameter requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. |

| Type | Restrictions |
|-----------------------------|--|
| | <ul style="list-style-type: none"> – Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. |
| Incremental synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • DDL operations performed on the source database will not be synchronized to the destination database. • The IMAGE, TEXT, and NTEXT big data types cannot be deleted. • You can add additional synchronization objects. |

Procedure

This section uses Microsoft SQL Server to Kafka synchronization as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 4-318 Synchronization task information

The screenshot shows a web form for creating a synchronization task. At the top, there is an orange warning banner 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 includes:

- 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 4-336 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 4-319 Synchronization instance details

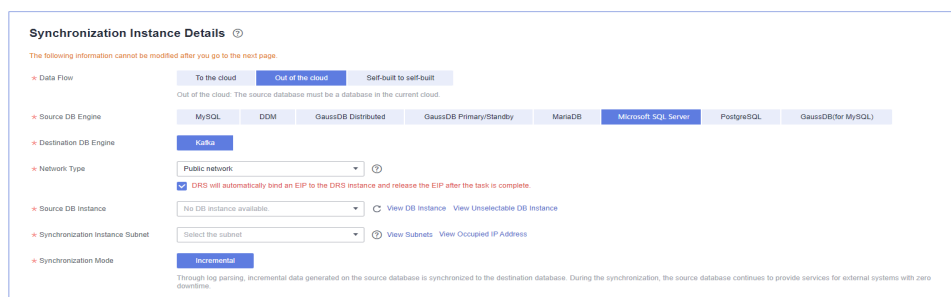


Table 4-337 Synchronization instance settings

| Parameter | Description |
|-----------------------|--------------------------------------|
| Data Flow | Select Out of the cloud . |
| Source DB Engine | Select Microsoft SQL Server . |
| Destination DB Engine | Select Kafka . |

| Parameter | Description |
|---------------------------------|--|
| Network Type | <p>Available options: VPC, Public network and VPN or Direct Connect. Public network is used as an example.</p> <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| Source DB Instance | Select an RDS for SQL Server instance you created as the source database. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Synchronization Mode | Incremental: Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- AZ

Figure 4-320 AZ




Table 4-338 Task AZ

| Parameter | Description |
|-----------|---|
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 4-321 Enterprise projects and tags

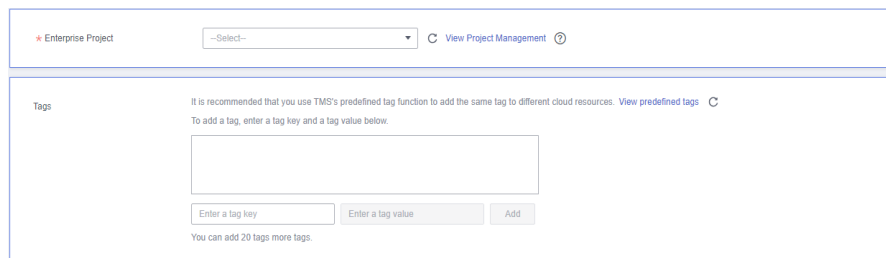


Table 4-339 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.


- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 4-322 Source database information

Source Database

DB Instance Name

Database Username

Database Password 

SSL Connection

✔ Test successful

Table 4-340 Source database settings

| Parameter | Description |
|-------------------------|---|
| DB Instance Name | The RDS for SQL Server instance you selected when creating the task. The parameter cannot be changed. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

Figure 4-323 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.


Security Protocol 

Table 4-341 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Security Protocol | Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication . |

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 4-324 Synchronization Object

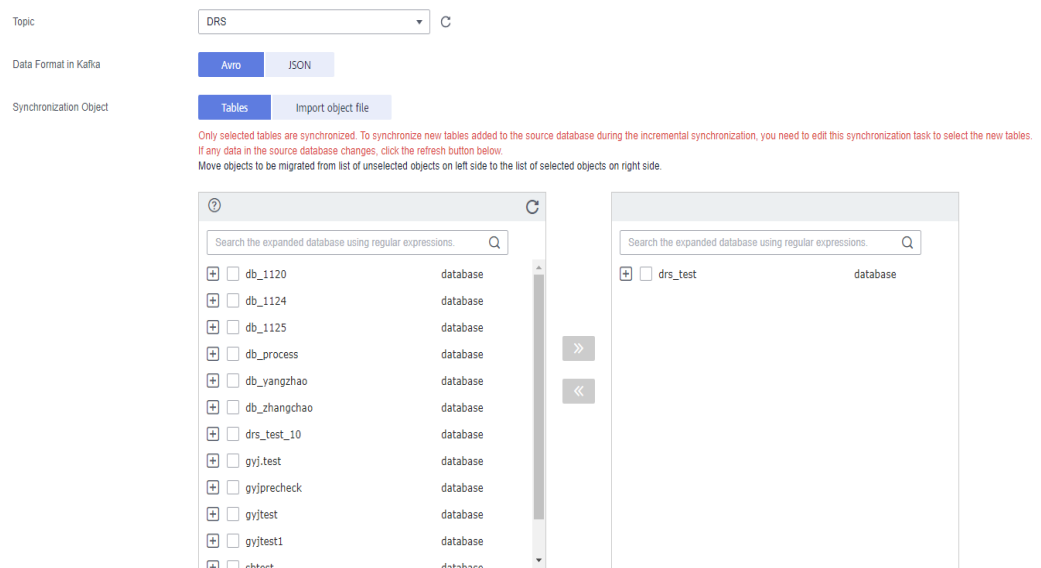



Table 4-342 Synchronization mode and object

| Parameter | Description |
|----------------------|--|
| Topic | Select the topic to be synchronized to the destination database. |
| Data Format in Kafka | Select the format of data delivered to Kafka. <ul style="list-style-type: none"> ● Avro refers to binary encoded format. ● Json refers to data interchange format. For details, see Kafka Message Format . |

| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. For details about the import procedure and description, Importing Synchronization Objects. When you select Import object file, you can use the mapping function in Mapping Object Names only when the topic synchronization policy is set to A specific topic. Otherwise, topics are generated based on the name format. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

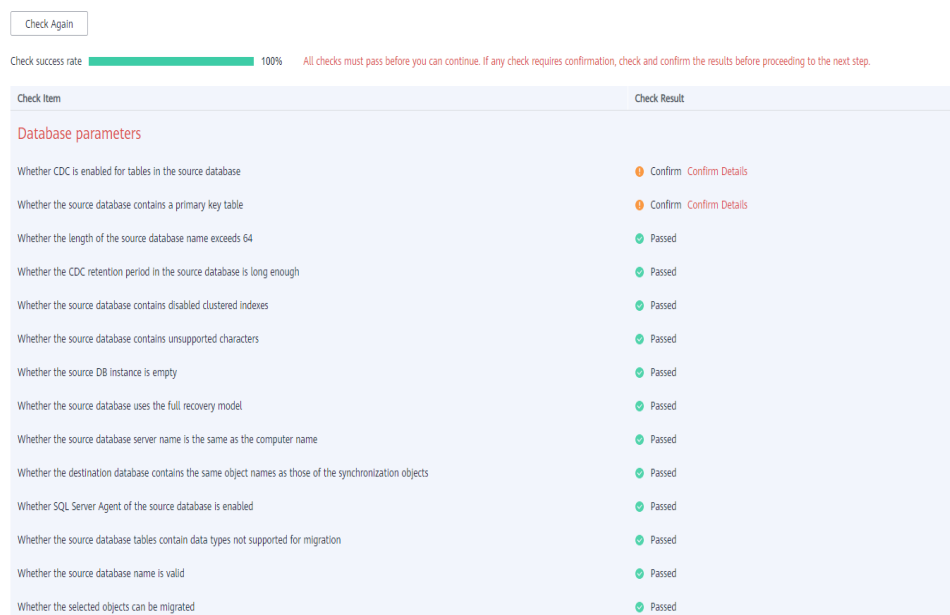
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 4-325 Pre-check



 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 4-326 Task startup settings

Table 4-343 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5 Between Self-built Databases

5.1 From MySQL to Kafka

Supported Source and Destination Databases

Table 5-1 Supported databases

| Source DB | Destination DB |
|---|---|
| <ul style="list-style-type: none">On-premises MySQL databasesMySQL databases on an ECS | <ul style="list-style-type: none">Kafka |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- Tables to be synchronized without a primary key may be locked for 3s.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full synchronization, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-2 Precautions

| Type | Restrictions |
|------------------------|---|
| Database permissions | <ul style="list-style-type: none"> • The source database user must have the following permissions: SELECT, LOCK TABLES, SHOW VIEW, EVENT, REPLICATION SLAVE, and REPLICATION CLIENT. |
| Synchronization object | <ul style="list-style-type: none"> • During full synchronization, tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized, but events and triggers cannot be synchronized. During incremental synchronization, only table data and DDLs can be synchronized. • Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|---|
| Source database | <ul style="list-style-type: none"> ● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format. ● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. <ul style="list-style-type: none"> - If the source database is an on-premises MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● The database and table names in the source database cannot contain non-ASCII characters, or 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. ● You are advised to set auto.create.topics.enable of Kafka to false. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. ● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● If a full synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the identifier field in the Kafka data for data deduplication. (The shard ID must be unique.) ● During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. ● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. ● If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● During the synchronization, do not delete or change the username, password, or permission of the source database, or change the port of the destination database. ● Data inconsistency may occur when the MyISAM table is modified during synchronization. ● During synchronization of table-level objects, renaming tables is not recommended. ● During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-1 Synchronization task information

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 5-3 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-2 Synchronization instance details

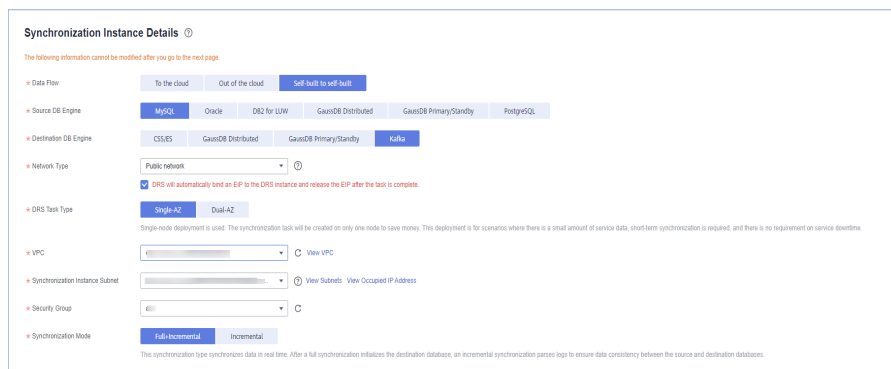


Table 5-4 Synchronization instance settings

| Parameter | Description |
|-----------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select MySQL . |
| Destination DB Engine | Select Kafka . |
| Network Type | The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect |
| DRS Task Type | Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none"> – Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. – Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task . |
| VPC | Select an available VPC. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . <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. |


- Task Type

Figure 5-3 Task type



Table 5-5 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |

| Parameter | Description |
|-----------|---|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-4 AZ</p>  |

- Enterprise Project and Tags

Figure 5-5 Enterprise projects and tags




Table 5-6 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-6 Source database information

Source Database

System databases, 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 5-7 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |


| Parameter | Description |
|----------------|--|
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-7 Destination database information

Destination Database

IP Address 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**.

The parameters on the **Set Synchronization Task** page vary according to the synchronization mode selected in **Step 2**. The **Full+Incremental** synchronization mode is used as an example in **Figure 5-8**.

Figure 5-8 Synchronization mode

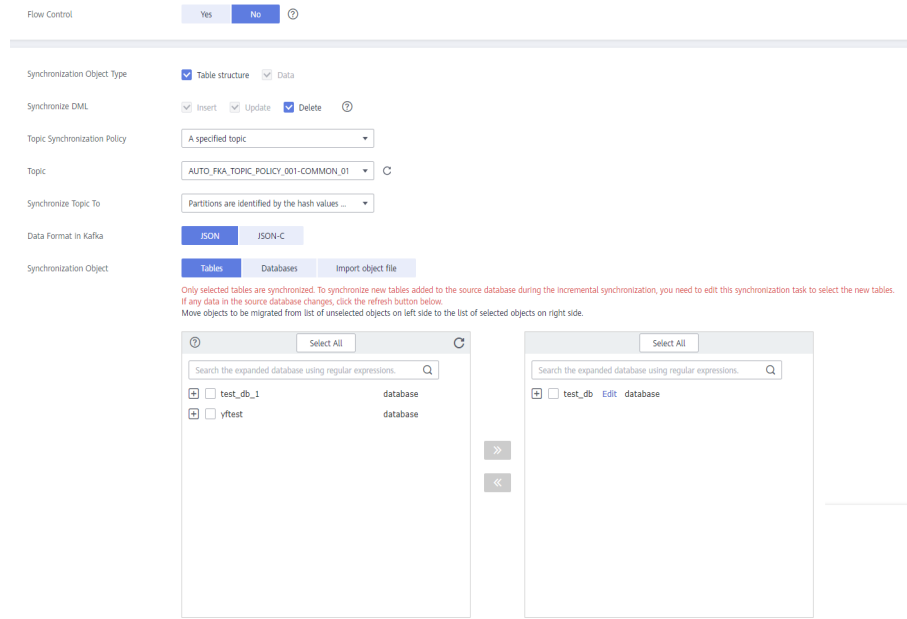
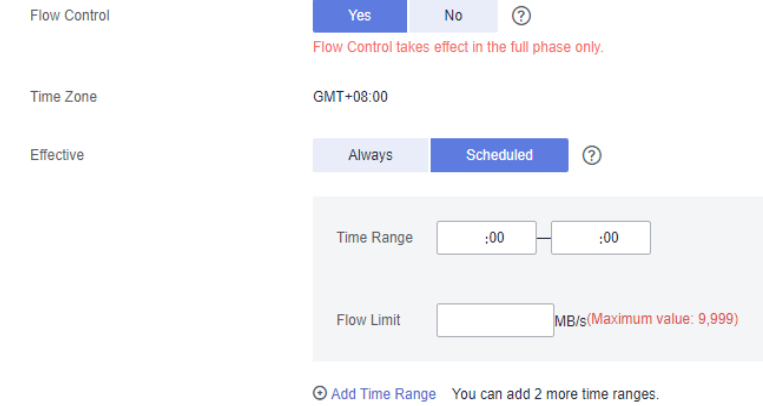



Table 5-9 Synchronization Object

| Parameter | Description |
|-----------------------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. |
| Synchronization Object Type | <p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> |
| Synchronize DML | <p>Select the DML operations to be synchronized. By default, all DML operations are selected.</p> |

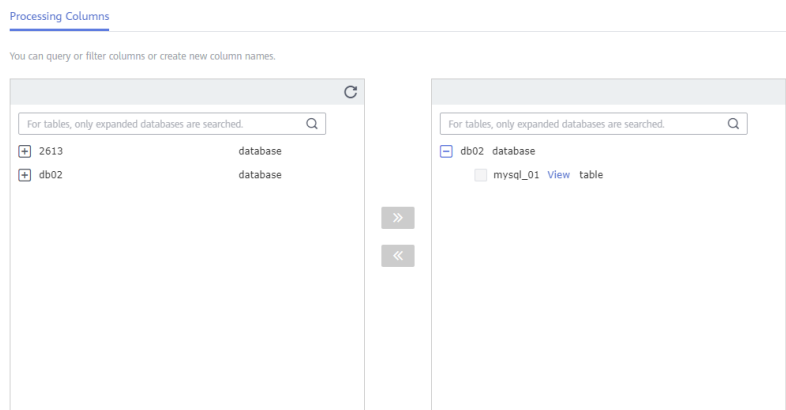
| Parameter | Description |
|------------------------------|---|
| Start Point | <p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| Topic Synchronization Policy | <p>Topic synchronization policy. You can select A specific topic or Auto-generated topics.</p> |
| Topic | <p>Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic.</p> |
| Topic Name Format | <p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>Only variables database and tablename are supported. The other characters must be constants. Replace \$database\$ with the database name and \$tablename\$ with the table name.</p> <p>For example, if this parameter is set to \$database\$-\$tablename\$ and the database name is db1, and the table name is tab1, the topic name is db1-tab1. If DDL statements are synchronized, \$tablename\$ is empty and the topic name is db1.</p> |
| Number of Partitions | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p> |
| Replication Factor | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p> |

| Parameter | Description |
|------------------------|--|
| Synchronize Topic To | <p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. • 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 from MySQL to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. This option is available only when Synchronization Mode is set to Incremental in Step 2. Only whitelisted users can use the Avro option. To use this option, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket. • JSON: JSON message format, which is easy to interpret but takes up more space. • JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

Figure 5-10 Processing data



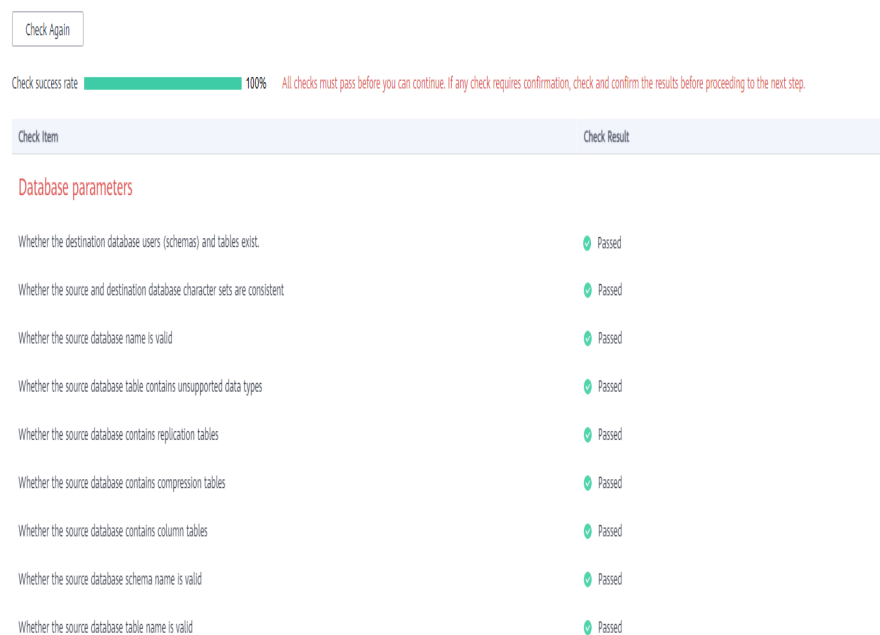
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Solutions to Failed Check Items](#) 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-12 Task startup settings


The screenshot shows the 'Task startup settings' interface. 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 toggle switch turned on; '* SMN Topic' is a dropdown menu; 'Delay Threshold (s)' is a toggle switch turned off; 'Data Exception Notification' is a toggle switch turned off; and '* Stop Abnormal Tasks After' is an input field containing the number '14', with a red note below it stating 'Abnormal tasks run longer than the period you set (unit: day) will automatically stop.'

Table 5-10 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----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 |

Suggestions

⚠ CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

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. • Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|-----------------|---|
| Source database | <ul style="list-style-type: none"> ● The source database names cannot contain non-ASCII characters, or the following characters: '<>\/\'' ● 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. <ul style="list-style-type: none"> - If the source database is an on-premises MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide. ● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● Enable skip-name-resolve for the MySQL source database to reduce the possibility of connection timeout. ● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The binary value is encrypted using Base64 and then written to the destination database. ● If no time zone is specified for the source database, specify the time zone for synchronizing the datetime type to the destination database. ● All table field names are converted to lowercase letters. ● If the <code>_id</code> field of the destination database is generated using multiple columns in the source database, separate these columns with colons (:). ● During task startup or full synchronization, you are not advised to perform DDL operations on the source database. ● To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL operations) during synchronization. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot write data using the statement-based binlog format. ● During the synchronization, do not clear binlogs on the source database. ● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. |

| Type | Restrictions |
|------|--|
| | <ul style="list-style-type: none"> • During the synchronization, do not create a database named ib_logfile in the source. • During incremental synchronization, if the source database is in a distributed transaction, the synchronization may fail. • Incremental synchronization filters out all DDL operations. • During incremental synchronization, resumable data transfer is supported. If the host system breaks down, data may be repeatedly inserted into non-transactional tables that do not have primary keys. • If table-level synchronization is selected, tables cannot be renamed during incremental synchronization. • Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • When you select synchronization objects, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the number of synchronization objects exceeds the limit, you can add synchronization objects in batches when you re-edit the synchronization objects. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-13 Synchronization task information

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, the form contains the following 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 (circle with 'i').
- Description:** A text area with a height of 0/256 characters and a help icon (circle with 'i').

Below the form, there is a note: "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 5-13 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-14 Synchronization instance details

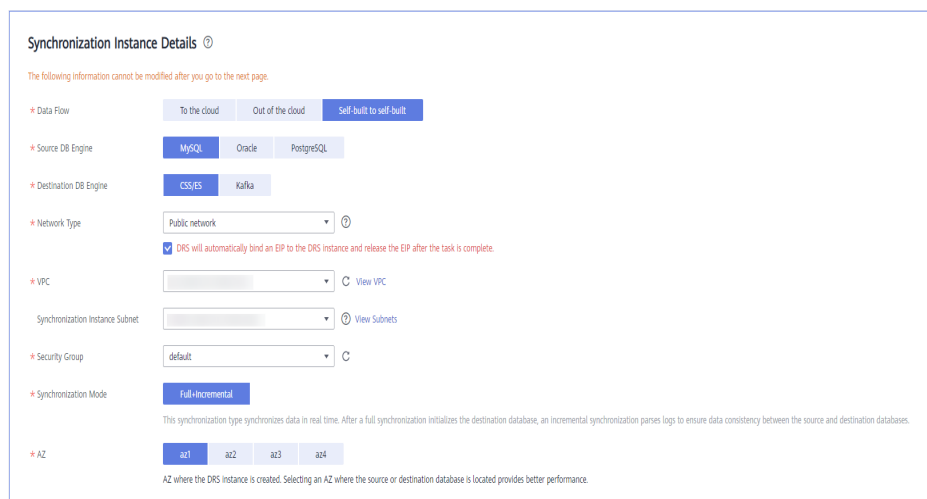


Table 5-14 Synchronization instance settings

| Parameter | Description |
|-----------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select MySQL . |
| Destination DB Engine | Select CSS/ES . |
| Network Type | The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental , data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. |

- Task Type

Figure 5-15 Task type



Table 5-15 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |

| Parameter | Description |
|-----------|---|
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-16 Enterprise projects and tags




Table 5-16 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |
| Tags | <ul style="list-style-type: none"> Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-17 Source database information

Source Database

System databases, 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 disabled, your data may be at risk. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-18 Destination database information

Destination Database

IP Address or Domain Name

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 disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 5-19 Synchronization mode

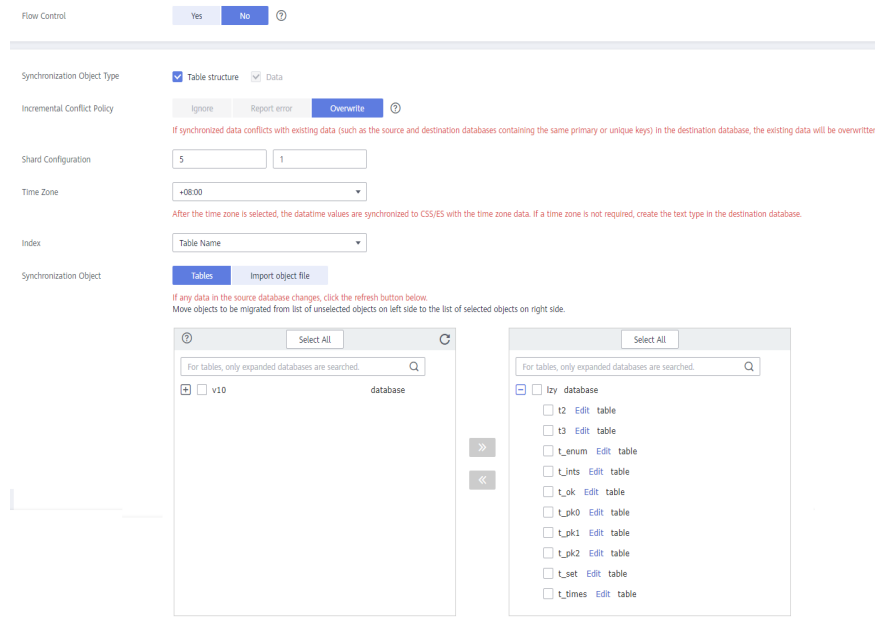
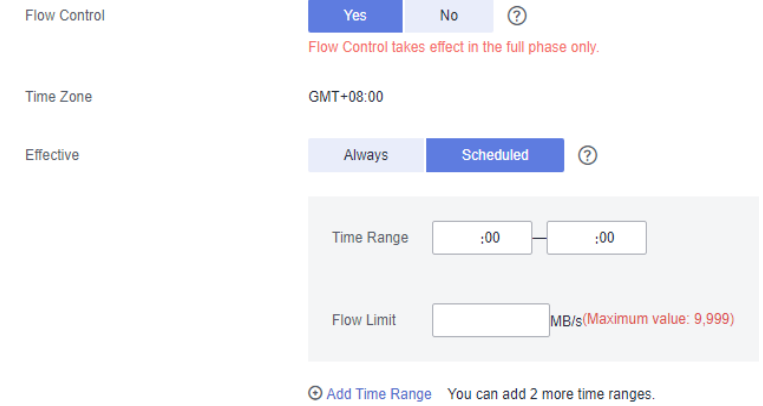



Table 5-19 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. |

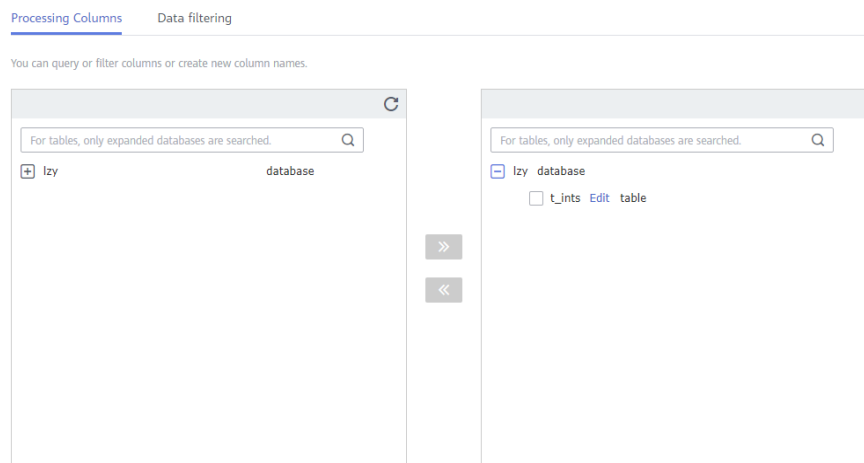
| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure or Data for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy only applies to incremental synchronization. The default value is Overwrite. The conflict in the full synchronization phase is ignored by default.</p> |
| Shard Configuration | <p>Configure the number of primary shards and the number of replicas. The default number of primary shards is 5, and the default number of shard replicas is 1.</p> |
| Time Zone | <p>After a time zone is selected, the datetime values are synchronized to CSS/ES with the time zone data.</p> |
| Index Name | <ul style="list-style-type: none"> • Table Name The index name created in the target Elasticsearch instance is the same as the table name. • Database Name+Table Name. The name of the index created in the target Elasticsearch instance is <i>DatabaseName_TableName</i>. |

| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> You can select objects by importing a file. For details, see Importing Synchronization Objects. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. If you select Tables, the size of the database table name and column name selected at a time cannot exceed 4 MB. If the limit is exceeded, you can add synchronization objects in batches by editing the synchronization objects. For details, see Editing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering** or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 5-21 Processing data



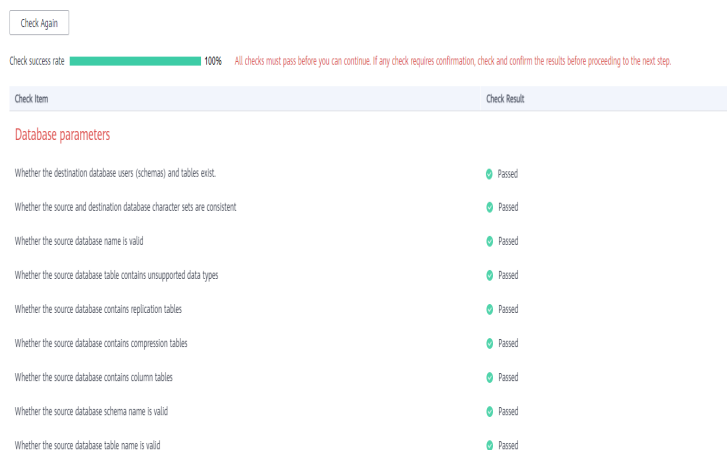
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-23 Task startup settings

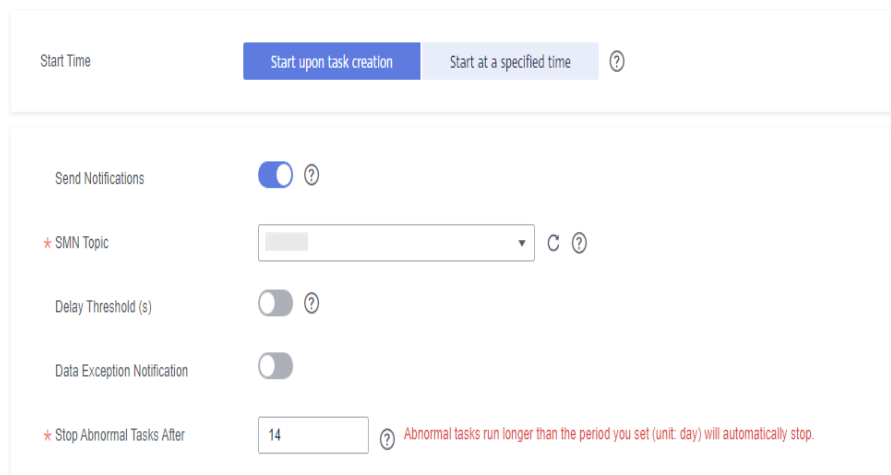



Table 5-20 Task startup settings

| Parameter | Description |
|-----------------------------|---|
| Start Time | <p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.</p> |
| Send Notifications | <p>This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification.</p> |
| SMN Topic | <p>This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber.</p> <p>For details, see Simple Message Notification User Guide.</p> |
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.3 From MySQL to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-21 Supported databases

| Source DB | Destination DB |
|--|-------------------------|
| <ul style="list-style-type: none"> • On-premises MySQL databases • MySQL databases on an ECS | GaussDB primary/standby |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-22 Precautions

| Type | Restrictions |
|----------------------|--|
| Database permissions | <ul style="list-style-type: none"> ● The source database user must have the following permissions: <ul style="list-style-type: none"> - Full synchronization: SELECT. Statement: GRANT SELECT ON <database>.<table> to <user>; - Full+incremental synchronization and incremental synchronization: SELECT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. Statement: GRANT SELECT, LOCK TABLES ON <database>.<table> to drsUser; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to <user>; ● 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. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> - Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> - Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● Only table structures, table data, and indexes can be synchronized. Other database objects such as stored procedures cannot be synchronized. ● Incremental synchronization does not support synchronization of DDL. ● Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed. ● The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and interval. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● During 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 a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. – If the source database is 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Restrictions |
|------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • The 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 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 table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • 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 an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated. • 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. |

| 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. • Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. • If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. • 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. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. • If 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 5-24 Synchronization task information




Table 5-23 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 5-25 Synchronization instance information




Table 5-24 Synchronization instance settings

| Parameter | Description |
|------------------|--|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select MySQL . |

| Parameter | Description |
|---------------------------------|--|
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | Public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 5-26 Task type



Table 5-25 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-27 Enterprise projects and tags




Table 5-26 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-28 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 5-27 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source database. |


| Parameter | Description |
|-------------------|--|
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.


Figure 5-29 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 


 Test successful

Table 5-28 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 synchronization objects and click **Next**.

Figure 5-30 Synchronization mode

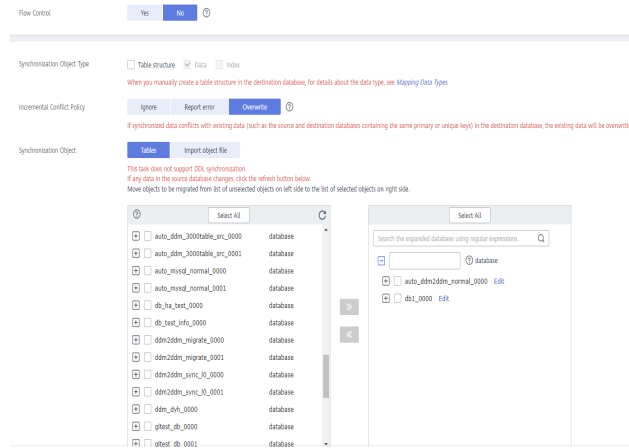
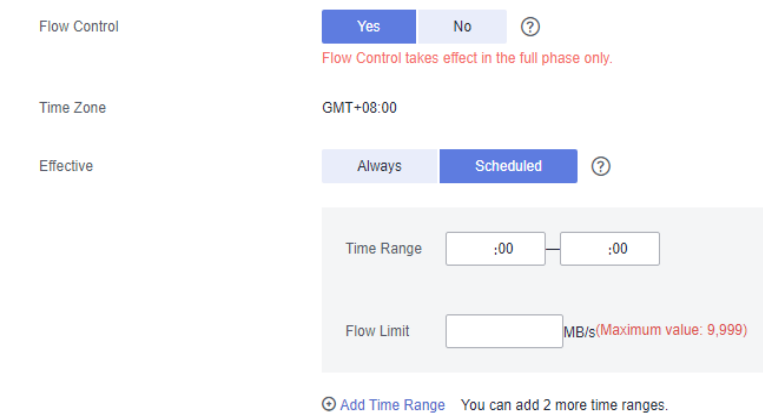



Table 5-29 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-31 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 5-32 Processing data

| Belonged Database/Table | New name | Column Name | Operation Type | Type | Operation |
|-------------------------|-------------|-------------|----------------|------|-----------|
| db1.jck_011 | db1.jck_011 | -- | -- | -- | Add |
| db1.tst2 | db1.tst2 | -- | -- | -- | Add |
| db1.tst3 | db1.tst3 | -- | -- | -- | Add |

Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-33 Pre-check

| Check Item | Check Result |
|---|-------------------------|
| Destination database storage space | |
| Whether the destination database has sufficient storage space | Passed |
| Database parameters | |
| Whether the source and destination database character sets are consistent | Confirm Confirm Details |
| Whether the source database contains tables without primary keys | Confirm Confirm Details |
| Whether the TIME_ZONE values of the source and destination databases are the same | Passed |
| Whether the SSL connection is correctly configured | Passed |
| Whether the source database binlog is row-based | Passed |
| Whether the binlog_row_image value of the source database is FULL | Passed |
| Whether the source database binlog is enabled | Passed |
| Whether objects with the same names exist in the source database | Passed |

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-34 Task startup settings





Table 5-30 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.4 From MySQL to GaussDB Distributed

Supported Source and Destination Databases

Table 5-31 Supported databases

| Source DB | Destination DB |
|--|---------------------|
| <ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS | GaussDB distributed |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-32 Precautions

| Type | Restrictions |
|----------------------|--|
| Database permissions | <ul style="list-style-type: none"> ● The source database user must have the following permissions: <ul style="list-style-type: none"> - Full synchronization: SELECT. Statement: GRANT SELECT ON <database>.<table> to <user>; - Full+incremental synchronization and incremental synchronization: SELECT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT. Statement: GRANT SELECT, LOCK TABLES ON <database>.<table> to drsUser; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to <user>; ● 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. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> - Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> - Table-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● Only table structures, table data, and indexes 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. ● Virtual columns are synchronized as common columns, and no data is written or incremental synchronization is performed. ● The following data types are not supported: XML, JSON containing the bit type, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and interval. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● 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. ● Full and incremental synchronizations do not support invisible columns. Invisible columns can be synchronized since MySQL 8.0.23. For example: <pre>CREATE TABLE `test11` (`id` int NOT NULL, `c1` int DEFAULT NULL /*!80023 INVISIBLE */, PRIMARY KEY (`id`));</pre> |

| Type | Restrictions |
|----------------------|--|
| Source database | <ul style="list-style-type: none"> ● During 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 a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. – If the source database is 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. ● During an incremental synchronization, if the session variable character_set_client is set to binary, some data may include garbled characters. ● 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. ● In MySQL, different tables in the same database (schema) can have the same index name or constraint name. In GaussDB, the index and constraint names are unique in the same schema, and the length is limited. To prevent conflicts between index and constraint names, the original index name in the table is changed to the following format after synchronization: hash value + original index name (which may be truncated) + _key. The hash value is calculated based on the original database name_original table name_original index name. Similarly, the original constraint name in the table is changed to the following format: hash value + original constraint name (which may be truncated) + _key. ● 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. ● Before creating a DRS task, if concurrency control rules of SQL statements are configured for the source database, the DRS task may fail. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> – Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Restrictions |
|------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. • The 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 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 table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. • 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 an incremental synchronization, tables whose primary key type is tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea, or binary cannot be deleted or updated. • 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. |

| 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. • Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario. • If a table without a primary key contains large fields (tinyblob, blob, mediumblob, longblob, tinytext, text, mediumtext, longtext, clob, nclob, bytea and binary), data of the large fields may be inconsistent during incremental synchronization. • Floating-point data is approximate numbers and depends on the OS platform and underlying implementation. FLOAT and DOUBLE data is inaccurate. If you synchronize floating-point data between MySQL and GaussDB, there may be some differences. For details, see the official MySQL documentation. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. • Two-phase commit 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 self-managed MySQL to GaussDB distributed as an example to describe how to configure a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-35 Synchronization task information




Table 5-33 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-36 Synchronization instance details

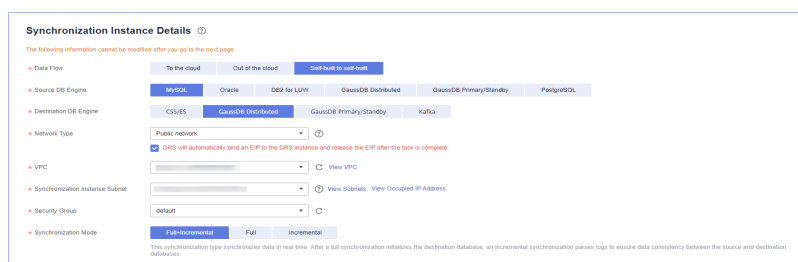


Table 5-34 Synchronization instance settings

| Parameter | Description |
|------------------|--|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select MySQL . |

| Parameter | Description |
|---------------------------------|--|
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization mode | The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles . <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 5-37 Task type



Table 5-35 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-38 Enterprise projects and tags




Table 5-36 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-39 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 5-37 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for accessing the source database. |


| Parameter | Description |
|-------------------|--|
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.


Figure 5-40 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 5-38 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 database username and password are encrypted and stored in the system and will be cleared after the task is deleted. |

Step 4 On the **Set Synchronization Task** page, select synchronization objects and click **Next**.

Figure 5-41 Synchronization mode

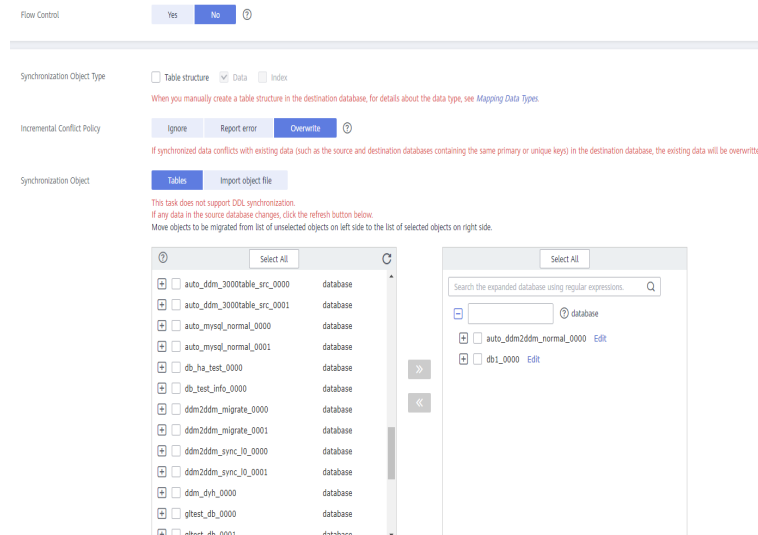
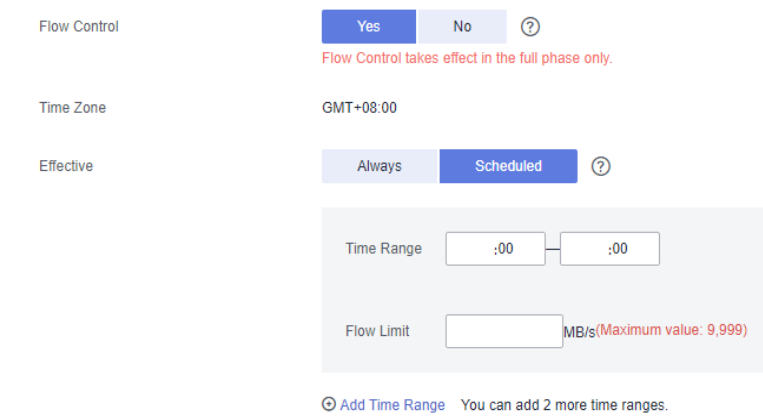



Table 5-39 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-42 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>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database. • Report error The synchronization task will be stopped and fail. • Overwrite Conflicting data will be overwritten. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

- Step 5** On the **Process Data** page, set the filtering rules for data processing.
- If data processing is not required, click **Next**.
 - If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 5-43 Processing data

| Selected Database/Table | Filter name | Column Name | Operation Type | Type | Operation |
|-------------------------|-------------|-------------|----------------|------|-----------|
| db1.pak_011 | db1.pak_011 | -- | -- | -- | Add |
| db1.test2 | db1.test2 | -- | -- | -- | Add |
| db1.test3 | db1.test3 | -- | -- | -- | Add |

- Step 6** On the **Check Task** page, check the synchronization task.
- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
 - If all check items are successful, click **Next**.

Figure 5-44 Pre-check

| Check Item | Check Result |
|---|-------------------------|
| Database parameters | |
| Whether there are source database tables that do not contain primary keys | Confirm Confirm Details |
| Whether the character set type is supported | Passed |
| Whether the destination database contains the configured databases | Passed |
| Whether the SSL connection is correctly configured | Passed |
| Whether the database names mapped to the destination database are valid | Passed |
| Whether the source database binlog is row-based | Passed |

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

- Step 7** On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-45 Task startup settings





Table 5-40 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.5 From Oracle to Kafka

Supported Source and Destination Databases

Table 5-41 Supported databases

| Source DB | Destination DB |
|--|---|
| <ul style="list-style-type: none"> On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) | <ul style="list-style-type: none"> Kafka |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-42 Environment Constraints

| Type | Constraint |
|----------------------|--|
| Database permissions | <ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. - Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. - To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs. |

| Type | Constraint |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● Only table data can be synchronized in real time. ● The following data types are supported: VARCHAR, VARCHAR2, NVARCHAR2, NUMBER, FLOAT, LONG, DATE, BINARY_FLOAT, BINARY_DOUBLE, CHAR, NCHAR, ROWID, TIMESTAMP, TIMESTAMP WITH TIME ZONE, and TIMESTAMP WITH LOCAL TIME ZONE. ● The following column types cannot pass the precheck: GEOMETRY and self-defined. ● The following column types cannot be synchronized but can pass the precheck: 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● Temporary tables in the source database cannot be synchronized. ● Tables whose default values contain expression functions cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. |

| Type | Constraint |
|----------------------|---|
| Source database | <ul style="list-style-type: none"> ● The names of databases and tables cannot contain non-ASCII characters or special characters .><`\ ,?!" ● An empty source database cannot be synchronized. ● If the source database is an RAC database, you cannot add or delete nodes. ● If the source database is an RAC database and uses SCAN IP, the synchronization instance must be able to connect to the virtual IP addresses of all RAC nodes. Otherwise, the connection check fails. ● Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15. |
| Destination database | <ul style="list-style-type: none"> ● The destination database is a Kafka database. ● You are advised to set auto.create.topics.enable of Kafka to false. |

| Type | Constraint |
|-------------|--|
| Precautions | <ul style="list-style-type: none"> ● If there are special characters such as Chinese and Japanese in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database. ● After data in the Oracle database is synchronized to Kafka, the character set becomes UTF8. ● The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record. ● For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience. ● If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address. ● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization. ● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency. ● If the PDB database is used for synchronization, all PDBs must be enabled during incremental synchronization due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● The supplemental log supports all or primary key+unique index columns. ● If a column that is not displayed in the log, it will not be displayed in the transferred message, which means that the column is not updated. ● During synchronization, do not delete the username, password, and permissions of the source and destination databases or modify the port of the destination database. |

| Type | Constraint |
|------|---|
| | <ul style="list-style-type: none"> • During the synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. • During synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails. • During the synchronization, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. • During incremental synchronization, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. • During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. • During incremental synchronization, do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • In an incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. • During an incremental synchronization of table-level objects, renaming tables is not recommended. • If you select Tables for Synchronization Object, all tables must be synchronized to the same topic at the destination end. If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. • DDL operations can be performed on tables. • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. |

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

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 [Info icon]

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 5-43 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-47 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**, PostgreSQL
- Destination DB Engine: **Kudu**
- Network Type: Public network [Info icon]
 - DRS will automatically bind an EIP to the DRS instance and release the EIP after the task is complete.
- DRS Task Type: **Single-AZ**, Dual-AZ

Single node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime.
- VPC: [Dropdown menu] [View VPC]
- Synchronization Instance Subnet: [Dropdown menu] [View Subnets] [View occupied IP address]
- Security Group: default [Info icon]
- 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 5-44 Synchronization instance settings

| Parameter | Description |
|---------------------------------|--|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select Oracle . |
| Destination DB Engine | Select Kafka . |
| Network Type | The Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect |
| DRS Task Type | Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none">- Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated.- Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task . |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <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. |

- Task Type

Figure 5-48 Task type



Table 5-45 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-49 AZ</p> <p>The screenshot shows two rows of AZ selection. The first row, labeled '* Primary AZ', has four tabs: az1 (selected), az2, az3, and az4. The second row, labeled '* Standby AZ', has four tabs: az1, az2, az3, and az4.</p> |

- Enterprise Project and Tags

Figure 5-50 Enterprise projects and tags

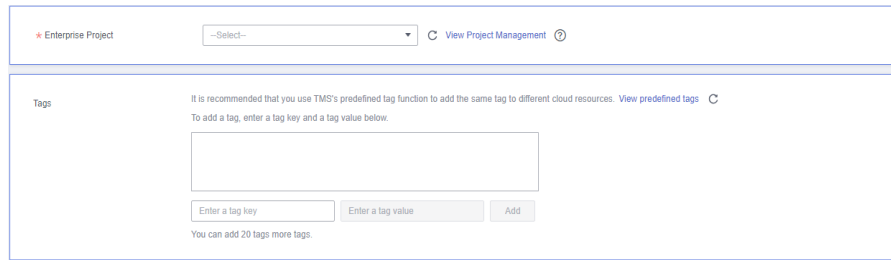


Table 5-46 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 5-51 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 5-47 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 disabled, your data may be at risk. |

 **NOTE**

The IP address, domain name, username, and password of the source database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Figure 5-52 Destination database information

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

Table 5-48 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 5-53 Synchronization mode

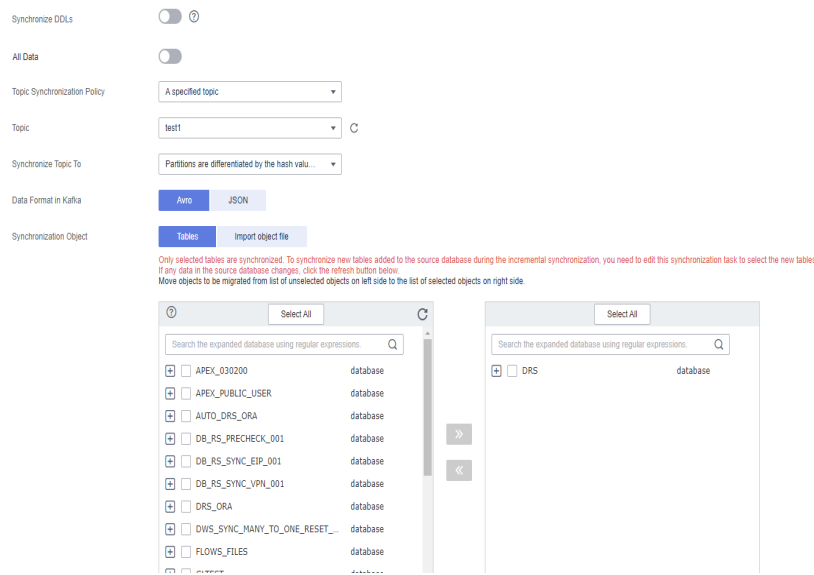



Table 5-49 Synchronization mode and object

| Parameter | Description |
|------------------------------|---|
| Synchronize DDLs | Controls whether to synchronize DDLs to Kafka. If Synchronize DDLs is enabled and Partitions are identified by the hash values of the primary key is selected, DDLs are hashed based on the table name because the DDLs do not have the primary key value. In other cases, the synchronization policy is the same as the partition policy. |
| All Data | Controls whether to synchronize all data in a single row. DRS parses the source database logs to synchronize incremental data. The data integrity in a single row depends on whether the values of all columns are recorded in the logs. If all data is required for the synchronization object, all-level supplemental logging must be enabled in the source database to record all column values of a single row data. This option is associated with the verification of the supplemental logging level in the source database in the pre-check phase. DRS incremental synchronization has the minimum requirement of table-level PK/UI supplemental logging. For details, see How Do I Check Supplemental Logging of the Source Oracle Database? |
| 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. |

| Parameter | Description |
|----------------------|---|
| Topic | <p>Select the topic to be synchronized to the destination database. This parameter is available when Topic Synchronization Policy is set to A specified topic.</p> |
| Topic Name Format | <p>Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> <p>The topic name format supports the schema and tablename variables. Other characters are used as constants. Replace \$schema\$ with the schema name and \$tablename\$ with the table name.</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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

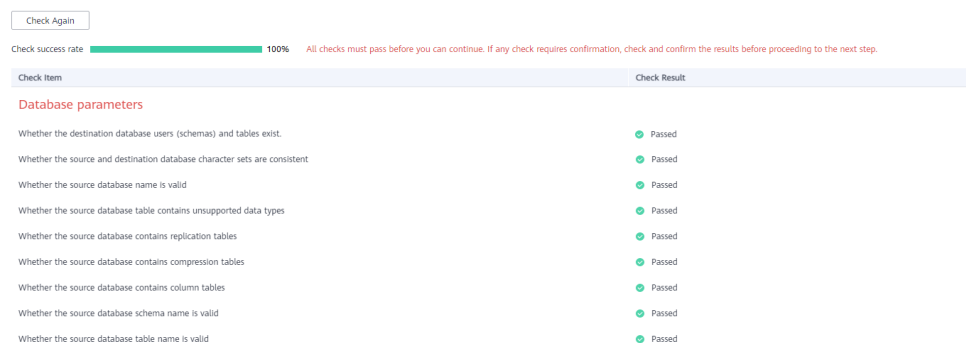
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-54 Pre-check



 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured

information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-55 Task startup settings

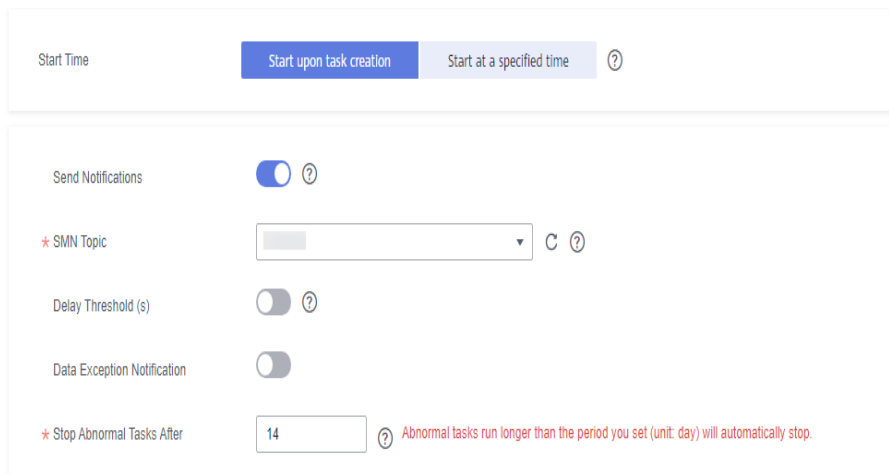



Table 5-50 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.6 From Oracle to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-51 Supported databases

| Source DB | Destination DB |
|---|-------------------------|
| <ul style="list-style-type: none">On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c)Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) | GaussDB primary/standby |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-52 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 ON <userName.tbName> to drsUser). - Full+incremental synchronization and incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"> • The destination database must have the following permissions: <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> - Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> - Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| 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. ● In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, and function indexes. ● For incremental synchronization, the LOB type supports only the BasicFiles attribute and does not support the SecureFiles attribute. The size of the LOB type must be less than 10 MB. ● For the TIMESTAMP WITH TIME ZONE data type, the 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 |

| Type | Restrictions |
|-----------------|--|
| | <p>3. Otherwise, incremental synchronization may fail because all columns cannot be matched.</p> <ul style="list-style-type: none"> ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● Tables with virtual columns in the source database cannot be synchronized. ● A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. ● When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. |
| 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 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 due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. ● 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. |

| Type | Restrictions |
|------|--|
| | <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. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● 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, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● Do not 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 |

| Type | Restrictions |
|------|--|
| | <p>synchronized, so you must manually create the table in the destination database.</p> <ul style="list-style-type: none"> ● During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. ● During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. ● In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● 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, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases |

| Type | Restrictions |
|------|--|
| | <p>requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database.</p> <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. • During incremental synchronization, 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 hidden columns (UNUSED, INVISIBLE). • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. • You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. • When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may occur after the switchover due to DDL changes during replication delay or long transactions. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database |

| Type | Restrictions |
|------|--|
| | sequence value minus the security margin. The default security margin is 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 5-56 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]

Table 5-53 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-57 Synchronization instance details

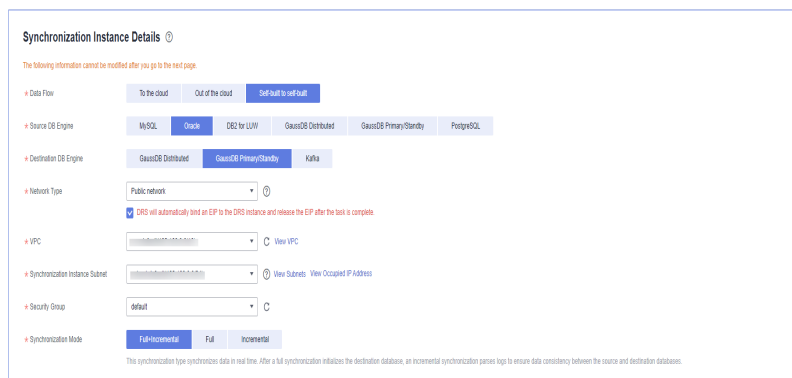


Table 5-54 Synchronization instance settings

| 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 Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | Available options: Full , Incremental , and Full +Incremental . |

- Task Type

Figure 5-58 Task type



Table 5-55 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-59 Enterprise projects and tags




Table 5-56 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-60 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

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 5-57 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance. |
| Port | The port of the source database. Range: 1 – 65535 |

| Parameter | Description |
|-----------------------|---|
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| PDB Name | Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

Figure 5-61 Destination database information

Destination Database

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 5-58 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 |

| Parameter | Description |
|-------------------|--|
| 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 5-62 Synchronization mode

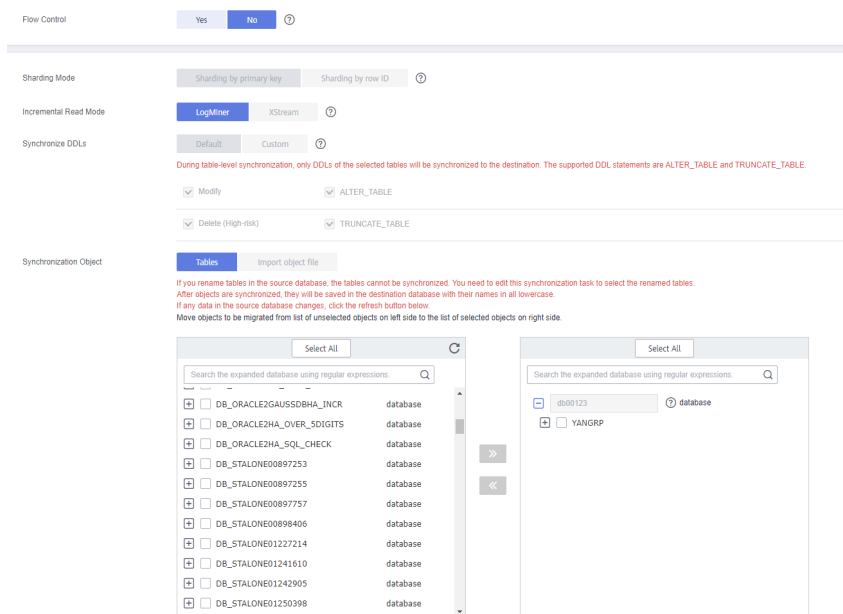
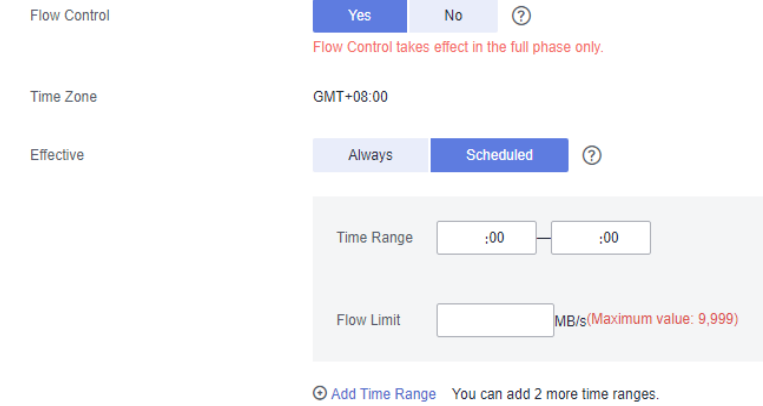



Table 5-59 Synchronization Object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-63 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

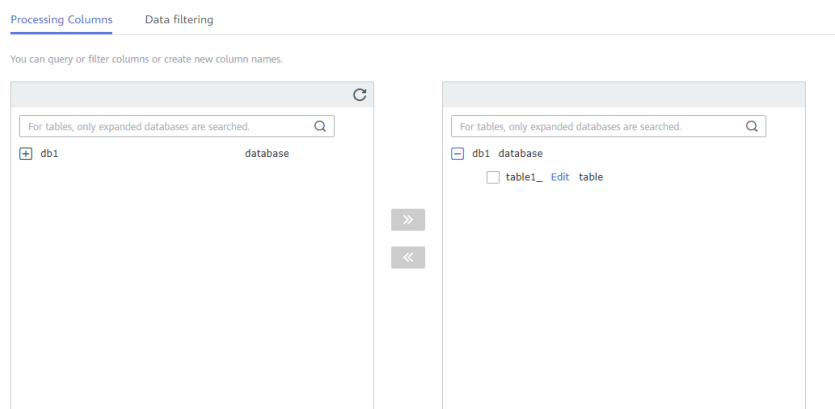
| Parameter | Description |
|-----------------------|--|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> • Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. • Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p> |
| Incremental Read Mode | <p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> • LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. • XStream: The XStream interface must be enabled. <p>LogMiner is recommended.</p> |
| Synchronize DDLs | <p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> • Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. • Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

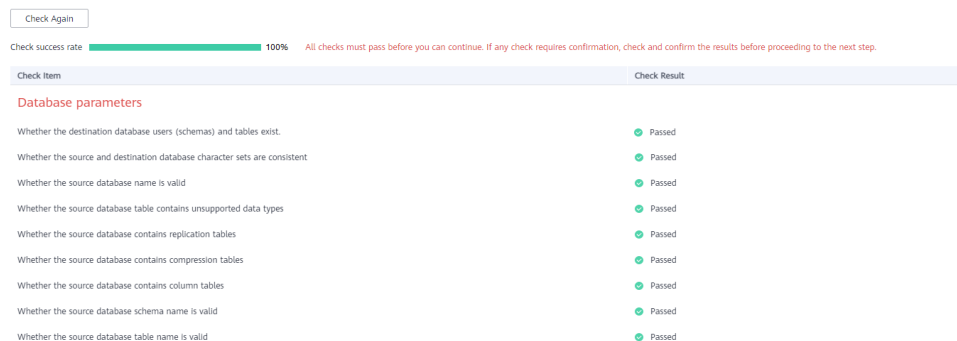
Figure 5-64 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.
- If all check items are successful, click **Next**.

Figure 5-65 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-66 Task startup settings





Table 5-60 Task startup settings

| Parameter | Description |
|------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |

| Parameter | Description |
|-----------------------------|--|
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.7 From Oracle to GaussDB Distributed

Supported Source and Destination Databases

Table 5-61 Supported databases

| Source DB | Destination DB |
|--|---------------------|
| <ul style="list-style-type: none"> On-premises databases (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c, 19c, and 21c) | GaussDB distributed |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full 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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-62 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 ON <userName.tbName> to drsUser). - Full+incremental synchronization and incremental synchronization: Oracle 12c or later in tenant mode: To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE and SELECT ANY TRANSACTION permissions for a PDB, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB. Oracle 12c or later in non-tenant mode: You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING. To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION. - During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled. - Oracle 12c or later does not support incremental synchronization using accounts whose |

| Type | Restrictions |
|------|--|
| | <p>ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.</p> <ul style="list-style-type: none"> • The destination database must have the following permissions: <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. Statement: GRANT CREATE, CONNECT ON DATABASE <database> TO <user> - Schema-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user> - Table-level permission: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the SCHEMA to the user. (The SELECT permission is required only for tables without primary keys.) Statement for granting the DML permission on all tables in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user> Statement for granting the DML permission on a specified table in the SCHEMA: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user> - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| 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. • In the incremental phase, Oracle extended characters are not supported. The standard character set cannot parse Oracle customized extended characters. • 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 full synchronization, the table structure does 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. |

| Type | Restrictions |
|-----------------|--|
| | <ul style="list-style-type: none"> • If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. • For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • Tables with virtual columns in the source database cannot be synchronized. • A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. • When you manually create a table structure, the time type in the destination database must be the same as that in the source database. Otherwise, time data may be inconsistent after time zone conversion. |
| 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 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 due to the restrictions of the Oracle LogMiner component. ● In Oracle 12.2 and later versions, due to the restrictions of the Oracle LogMiner component, a schema, table, or column name contains no more than 30 characters during an incremental synchronization. ● If the total length of the storage data in each row exceeds the upper limit (8192 bytes by default) of GaussDB, the task may fail. ● 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. |

| Type | Restrictions |
|------|--|
| | <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. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is ZHS16GBK and the destination character set is UTF8, a Chinese character of ZHS16GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases. ● 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, the username (schema name) of the source Oracle database cannot be changed, including the scenarios where the schema name is changed by modifying the USERS\$ dictionary table in versions earlier than 11.2.0.2 and by using ALTER USER username RENAME TO new_username in versions later than 11.2.0.2. ● Do not 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 |

| Type | Restrictions |
|------|--|
| | <p>synchronized, so you must manually create the table in the destination database.</p> <ul style="list-style-type: none"> ● In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. ● During table structure synchronization in a full synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created. ● During table structure synchronization in a full synchronization, the column name cannot be CTID, XMIN, CMIN, XMAX, CMAX, TABLEOID, XC_NODE_ID, TID, OID, GS_TUPLE_UID or TABLEBUCKETID. Otherwise, the table fails to be created. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. During incremental synchronization, operations on partitioned tables in the source database may fail to be synchronized to the destination database or fail to be executed in the destination database. ● 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, you are not advised to select a hybrid partition table because DML logs are not generated when data in the external partition of the hybrid partition table changes. DRS cannot obtain the changes during incremental synchronization, which may cause data inconsistency. ● During incremental synchronization, the incremental parsing supports only data within the valid range of the Oracle time type. Data beyond the valid range may cause task exceptions. For example, data truncation occurs if the year value is greater than 9999. ● In a full+incremental or incremental synchronization, the PDB database cannot be directly connected. You need to provide the service name/SID, username, and password of the CDB. ● During incremental synchronization, some DDL operations are supported. DDL conversion of heterogeneous databases |

| Type | Restrictions |
|------|--|
| | <p>requires semantic analysis and syntax compatibility. Only some DDL operations can be synchronized when the conversion is successful and the following conditions are met. If the conversion or execution fails, the error will be ignored. You need to manually execute the corresponding DDLs in the destination database.</p> <ul style="list-style-type: none"> - Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported. - The character length of new fields in incremental DDLs is not automatically extended. • During incremental synchronization, 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 hidden columns (UNUSED, INVISIBLE). • 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. • You are advised to disable the global secondary index (GSI) function for the destination database. Otherwise, incremental synchronization may fail. • When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours. • You are not advised to use the LOB type and extended character type (the length exceeds 4000 bytes) as incremental data filtering conditions. Oracle logs may not record the old value of update. • When a dual-AZ task is switched over (for example, the subtask in the primary AZ fails), if the primary task has a long delay or there are transactions that have been there for a long time without being submitted in the source database, the task may fail to be started after the switchover because the source database logs are cleared, or synchronization exceptions may |

| Type | Restrictions |
|------|--|
| | <p>occur after the switchover due to DDL changes during replication delay or long transactions.</p> <ul style="list-style-type: none"> If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. |

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

Table 5-63 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-68 Synchronization instance details

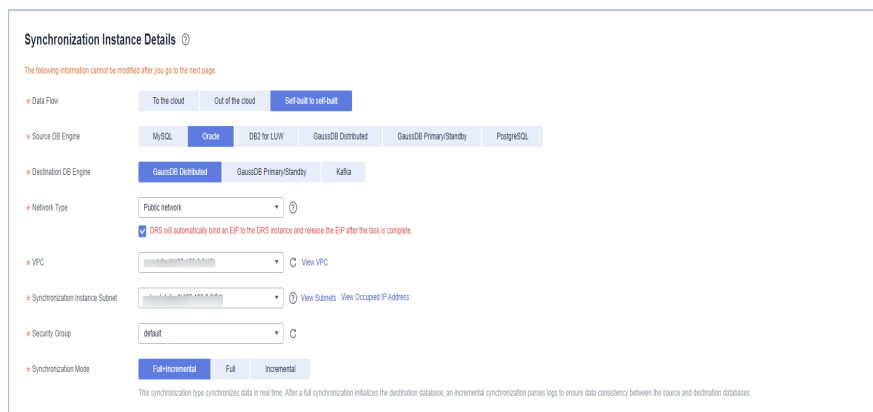


Table 5-64 Synchronization instance settings

| 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 Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | Available options: Full , Incremental , and Full +Incremental . |

- Task Type

Figure 5-69 Task type



Table 5-65 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-70 Enterprise projects and tags

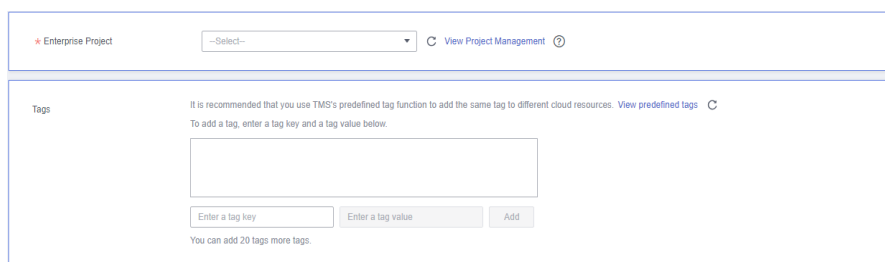


Table 5-66 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-71 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name
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 5-67 Source database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| PDB Name | Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

| Parameter | Description |
|----------------|--|
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

Figure 5-72 Destination database information

Destination Database

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

This button is available only after the replication instance is created successfully.

Table 5-68 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 5-73 Synchronization mode

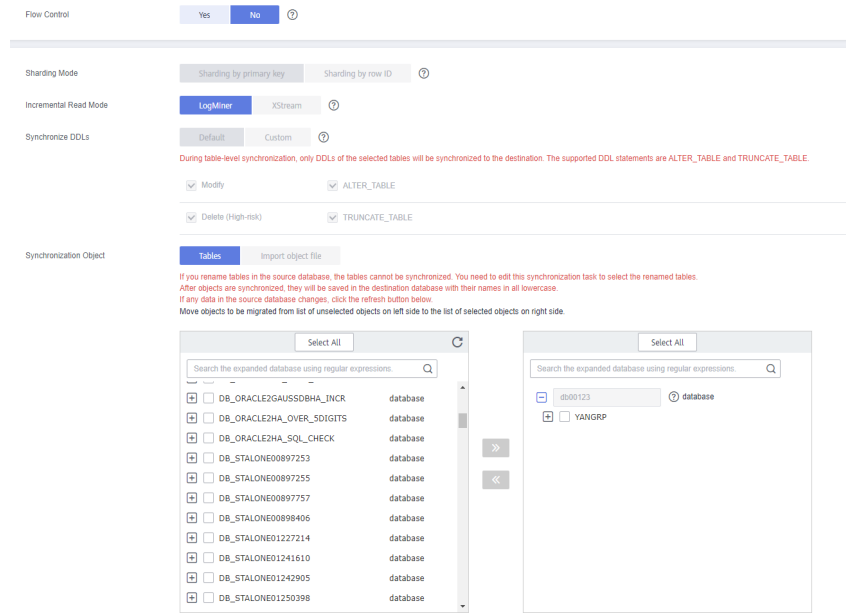
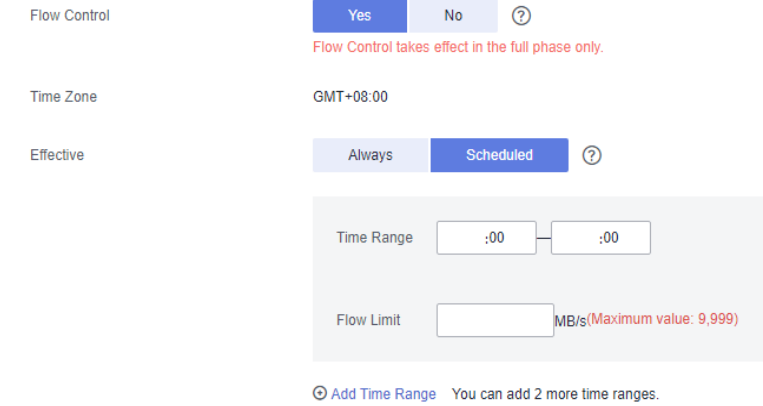



Table 5-69 Synchronization Object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-74 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

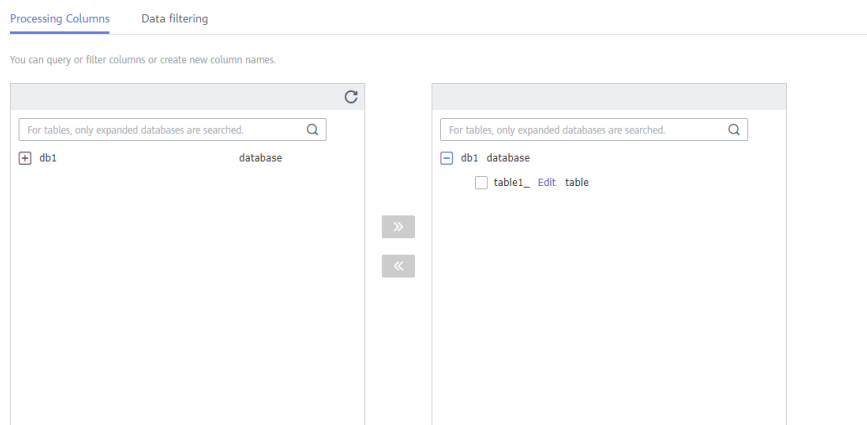
| Parameter | Description |
|-----------------------|--|
| Sharding Mode | <p>Select Sharding by primary key or Sharding by row ID as needed.</p> <ul style="list-style-type: none"> ● Sharding by primary key Tables with primary keys are sharded using primary key values. The sharding efficiency is lower than that of Sharding by row ID. ● Sharding by row ID Sharding by row ID is recommended for synchronizing large tables without primary keys, improving sharding efficiency. <p>Restrictions on Sharding by row ID: During full synchronization, do not perform operations that will cause row ID changes on the source database, such as Export/import of the table, ALTER TABLE XXXX MOVE, ALTER TABLE XXXX SHRINK SPACE, FLASHBACK TABLE XXXX, Splitting a partition, Updating a value so that it moves to a new partition, or Combining two partitions. Otherwise, data may be inconsistent. You can run ALTER TABLE XXXX DISABLE ROW MOVEMENT to disable these operations.</p> |
| Incremental Read Mode | <p>Select the mode of reading logs in the incremental synchronization phase.</p> <ul style="list-style-type: none"> ● LogMiner: uses the official Oracle interface to read redo logs. This mode is stable. ● XStream: The XStream interface must be enabled. <p>LogMiner is recommended.</p> |
| Synchronize DDLs | <p>DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default During table-level synchronization, only DDLs of the selected tables will be synchronized to the destination. The supported DDL statements are CREATE_TABLE, ALTER_TABLE, and TRUNCATE_TABLE. ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-75 Processing data



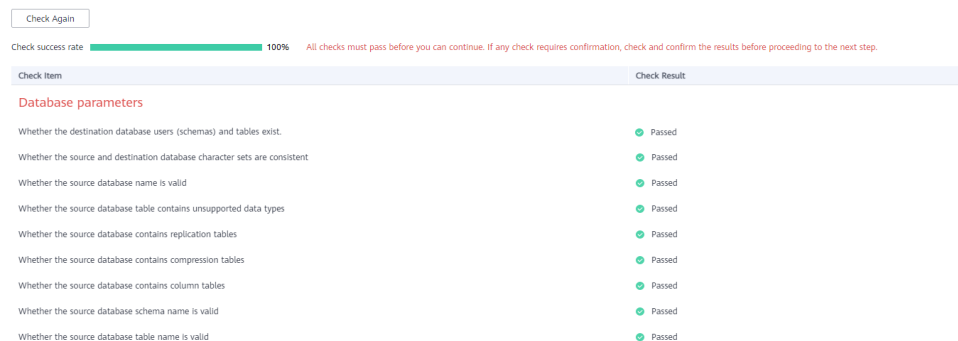
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-76 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 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-77 Task startup settings





Table 5-70 Task startup settings

| Parameter | Description |
|------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |

| Parameter | Description |
|-----------------------------|--|
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none">• If the delay threshold is set to 0, no notifications will be sent to the recipient.• In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent.• Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none">• You can set this parameter only for pay-per-use tasks.• Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.8 From PostgreSQL to Kafka

Supported Source and Destination Databases

Table 5-71 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 5-72 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-72 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 5-73](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-73 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 5-74 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> The wal_level value of the source database must be logical. The test_decoding plug-in has been installed on the source database. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the synchronization task cannot be created. - During database-level synchronization, incremental synchronization tasks cannot be edited. - During table-level synchronization, incremental synchronization tasks can be edited, but the database cannot be changed. - Before starting a synchronization task, ensure that no long transaction is started in the source database. Otherwise, the creation of the logical replication slot will be blocked, leading the task to fail. - After a task is started, a primary/standby switchover can be performed only on the source database of the following versions: RDS for PostgreSQL 12.6 or later and RDS for PostgreSQL 13 or later. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. |

| Type | Constraints |
|-----------------------------|---|
| Incremental synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not delete the primary key of the source database table. Otherwise, incremental data may be lost or the task may fail. • Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be lost or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, incremental data may be lost or the task may fail. • During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail. • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. |
| Stopping a task | <ul style="list-style-type: none"> • Stop a task normally: <ul style="list-style-type: none"> - When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted. • Forcibly stop a task: <ul style="list-style-type: none"> - To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL. - The naming rule of a logic replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.

- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-78 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 5-75 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-79 Synchronization instance details

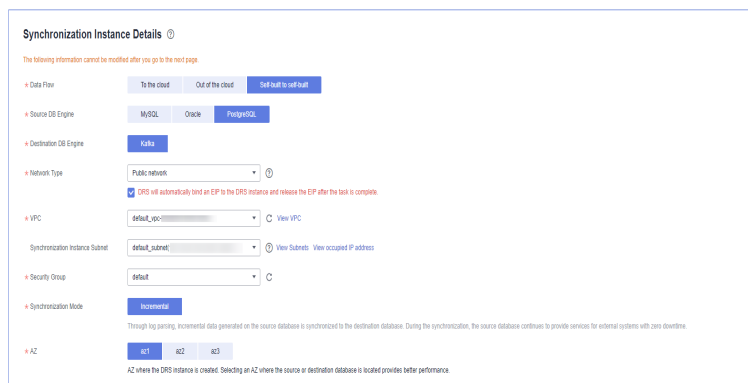


Table 5-76 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select PostgreSQL . |
| Destination DB Engine | Select Kafka . |
| Network Type | The public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | – 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. |

- Task Type

Figure 5-80 Task type



Table 5-77 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-81 Enterprise projects and tags

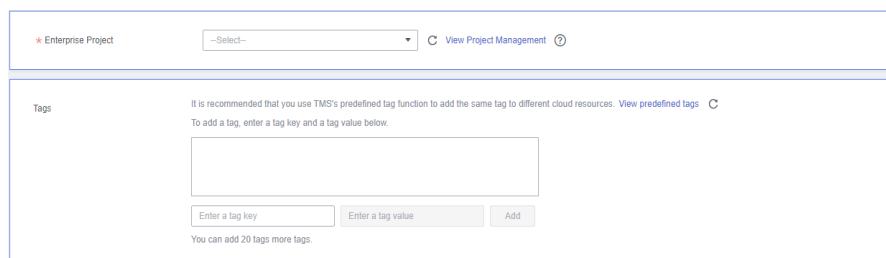


Table 5-78 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-82 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 5-79 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 5-83 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 5-80 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 5-84 Synchronization mode

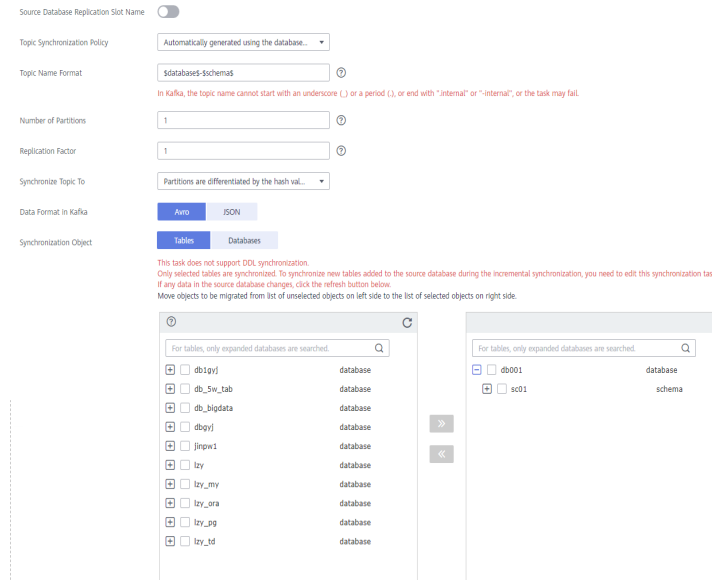



Table 5-81 Synchronization object

| 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 <code>database_name.schema.table_name</code>. ● 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

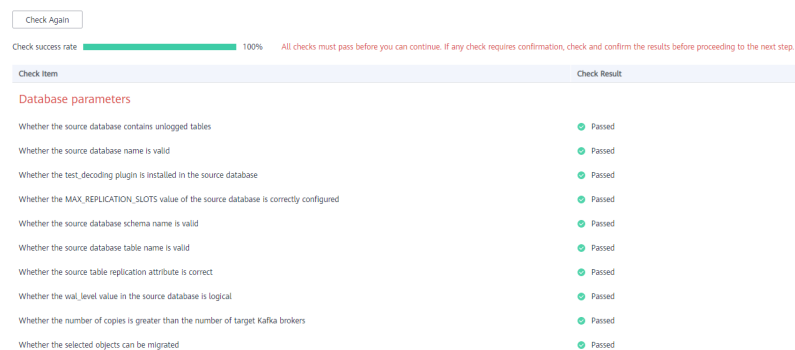
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-85 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-86 Task startup settings





Table 5-82 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.9 From GaussDB Primary/Standby to MySQL

Supported Source and Destination Databases

Table 5-83 Supported databases

| Source DB | Destination DB |
|-------------------------|--|
| GaussDB primary/standby | <ul style="list-style-type: none"> On-premises MySQL 5.5, 5.6, and 5.7 databases ECS-hosted MySQL 5.5, 5.6, and 5.7 databases MySQL 5.5, 5.6, and 5.7 databases on other clouds |

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 5-84 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-84 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. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 5-85](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-85 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|--|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 5-86 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 names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. - In GaussDB, if there is 0x00 in character or binary data, the data will be truncated, which is inconsistent with MySQL. As a result, the final data may be inconsistent. - Binary data cannot be used as a DML matching condition. That is, binary data cannot be used as a primary key. Binary data is fixed in length and is automatically filled in. The fill character of GaussDB is 0x20, and that of MySQL is 0x00. |
| 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 GaussDB Developer Guide. |
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Do not limit the synchronization speed during data comparison. • MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. |

| Type | Restrictions |
|-----------------|---|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 5-87 Synchronization task information

Table 5-87 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-88 Synchronization instance details

Table 5-88 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 MySQL . |
| Network Type | Public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| 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. |

- Task Type

Figure 5-89 Task type



Table 5-89 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-90 Enterprise projects and tags




Table 5-90 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.


Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.


Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.


Figure 5-91 Source database information

Source Database

DR Synchronization 

IP Address or Domain Name 
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-91 Source database settings

| Parameter | Description |
|---------------------------|--|
| DR Synchronization | If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name . |
| IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.


Figure 5-92 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password 

SSL Connection


 Test successful

Table 5-92 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 |

| 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 CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-93 Synchronization mode

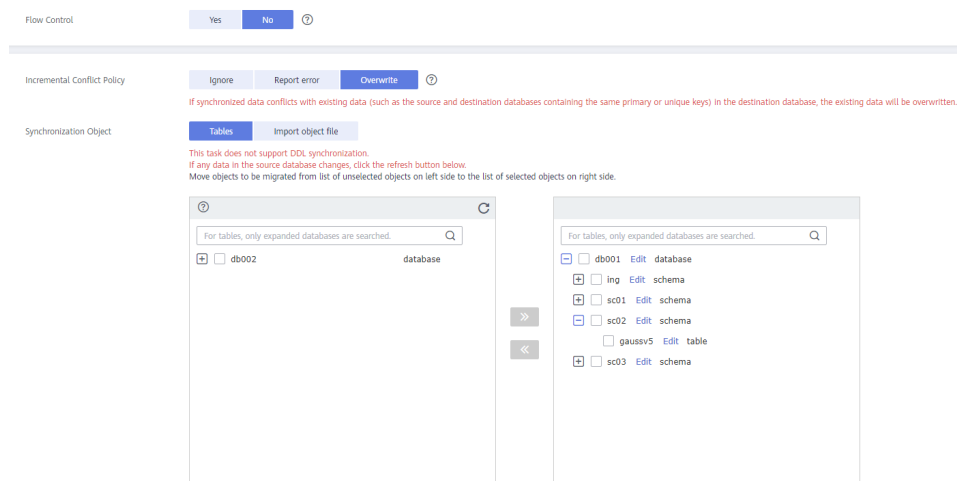
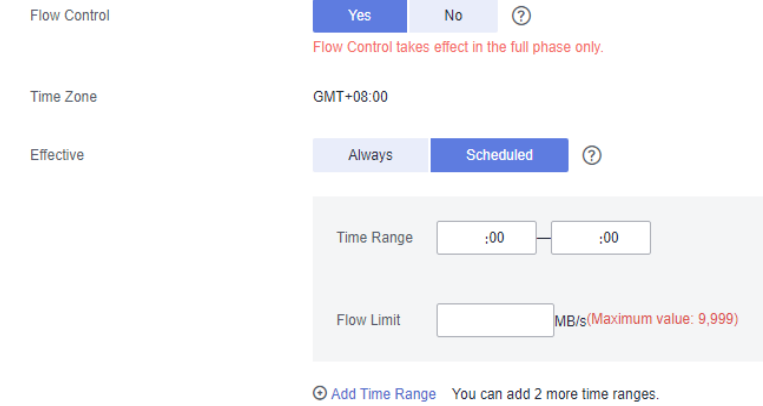



Table 5-93 Synchronization object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. |

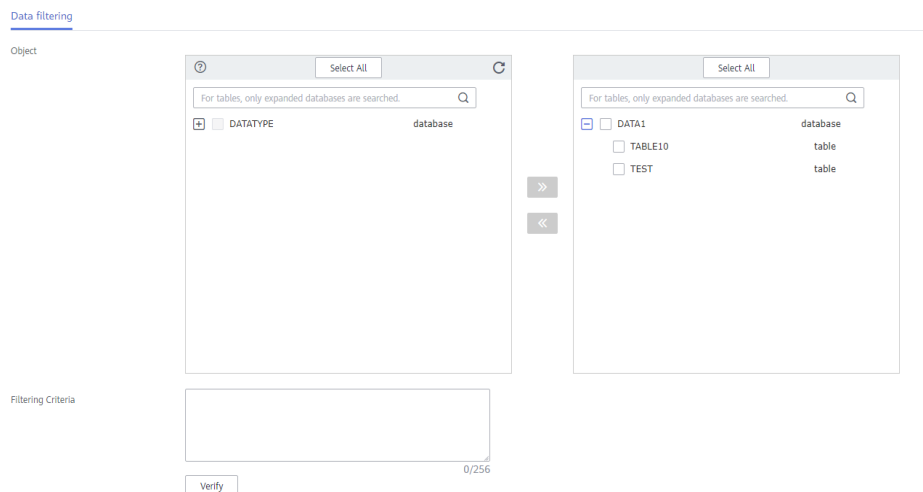
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-95 Processing data



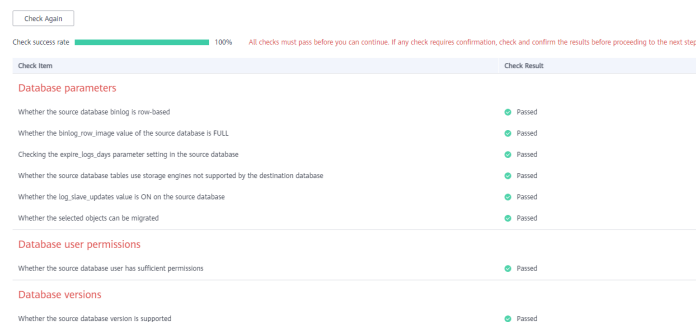
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-96 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-97 Task startup settings





Table 5-94 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.10 From GaussDB Primary/Standby to Oracle

Supported Source and Destination Databases

Table 5-95 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 5-96](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-96 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 and some DDL statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - 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 5-97](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-97 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|---|--|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT and INSERT permissions for tables, or the RESOURCE role. | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role | The user must have the DBA role or the following minimum permissions: CREATE SESSION, SELECT, INSERT, UPDATE, and DELETE permissions for tables or the RESOURCE role |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 5-98 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 REPLICIA 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. - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails. - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task |

| Type | Restrictions |
|----------------------|--|
| | <p>failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs.</p> <ul style="list-style-type: none"> - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding of GaussDB Developer Guide". ● 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. ● The following types of DDL statements cannot be synchronized: ALTER TABLE MERGE PARTITIONS and CREATE INDEX CONCURRENTLY. ● To synchronize DDL statements, ensure that the compatibility modes of the source and destination databases are the same. ● For DDL statements involving tablespace operations, the user of the destination database must have the tablespace operation permissions. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

| Type | Restrictions |
|-----------------|---|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 5-98 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 [Info icon]

Description: [Text area] [Info 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 5-99 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-99 Synchronization instance details

Synchronization Instance Details

The following information cannot be modified after you go to the next page.

- Data Flow: To the cloud, Out of the cloud, **Self-built to self-built**
- Source DB Engine: MySQL, Oracle, DB2 for LUW, GaussDB Distributed, **GaussDB Primary/Standby**, PostgreSQL
- Destination DB Engine: Oracle, **GaussDB DWS**, GaussDB Primary/Standby, Kafka
- Network Type: Public network [Info icon]
- Network Type Note: DRS will automatically bind an EIP to the DRS instance and release the EIP after the task is complete.
- VPC: [Dropdown menu] [View VPC]
- Synchronization Instance Subnet: [Dropdown menu] [View Subnets] [View Occupied IP Address]
- Security Group: default [C]
- 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 passes logs to ensure data consistency between the source and destination databases.

Table 5-100 Synchronization instance settings

| Parameter | Description |
|-----------|--|
| Data Flow | Choose Self-built to self-built . |

| Parameter | Description |
|---------------------------------|---|
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select Oracle . |
| Network Type | The public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 5-100 Task type



Table 5-101 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-101 Enterprise projects and tags




Table 5-102 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.


Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.


Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.


Figure 5-102 Source database information

Source Database

DR Synchronization 

IP Address or Domain Name 
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-103 Source database settings

| Parameter | Description |
|---------------------------|--|
| DR Synchronization | If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name . |
| IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-103 Destination database information

Destination Database


IP Address or Domain Name

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 5-104 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-104 Synchronization mode

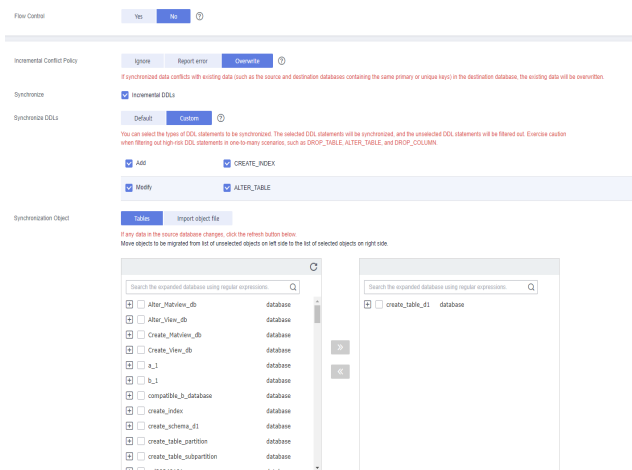
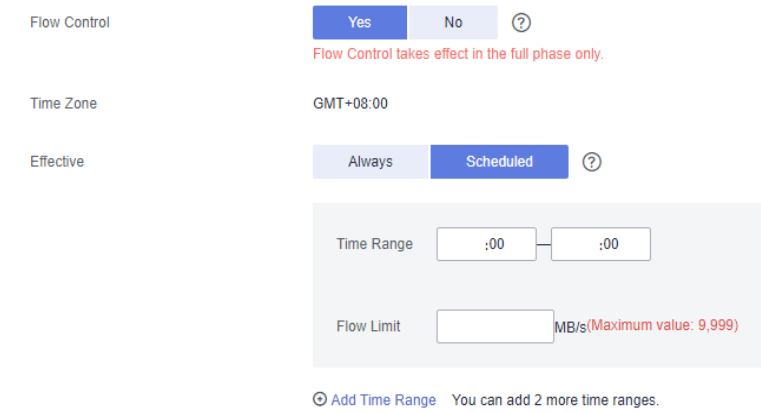



Table 5-105 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-105 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> - The flow control mode takes effect only in the full synchronization phase. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

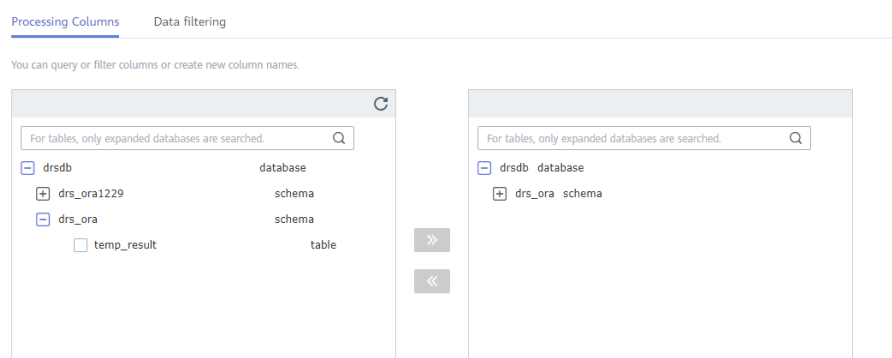
| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| Synchronize | <p>Indicates whether to synchronize DDLs in the incremental synchronization phase.</p> |
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: none"> ● Default: DDL statements of the selected synchronization objects will be synchronized to the destination database. <ul style="list-style-type: none"> – DDL statements supported by table-level synchronization: ALTER TABLE and CREATE INDEX ● Custom: You can select the DDL type to be synchronized as required. The DDL types supported by different data flow types are displayed on the GUI. Only selected DDL types can be synchronized. |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object files for Synchronization Object based on your service requirements.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-106 Processing data



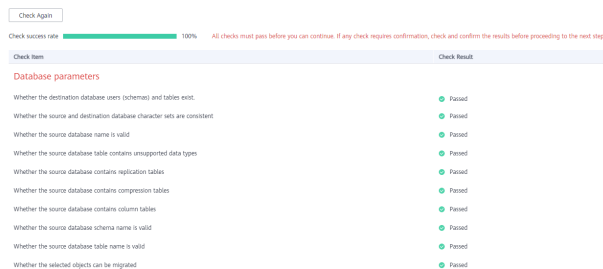
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-107 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-108 Task startup settings

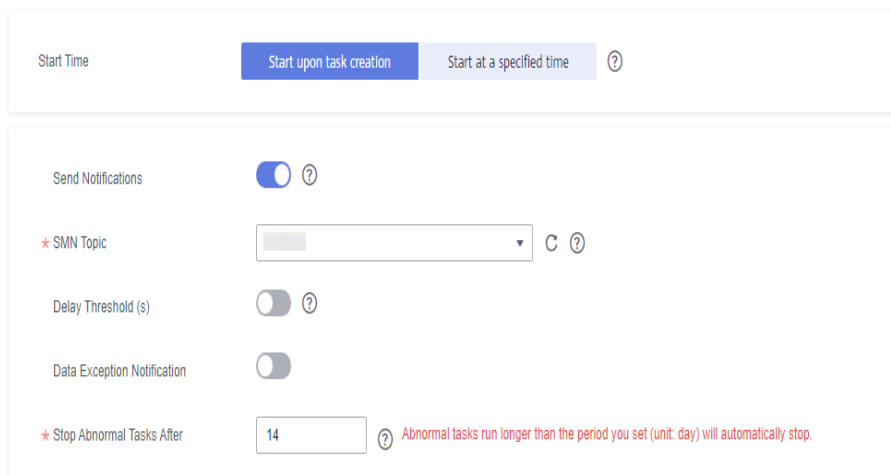



Table 5-106 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |

| Parameter | Description |
|-----------------------------|--|
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.11 From GaussDB Primary/Standby to Kafka

Supported Source and Destination Databases

Table 5-107 Supported databases

| Source DB | Destination DB |
|-------------------------|---------------------|
| GaussDB primary/standby | Kafka 0.11 or later |

Supported Synchronization Objects

Table 5-108 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-108 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 • The supported fields are BIGINT, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE and TINYINT. • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of selected tables can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \"?!. The column name cannot contain double quotation marks (\"), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 5-109](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-109 Database user permission

| Type | Incremental |
|----------------------|---|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. • The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. • Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 5-110 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?!. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - If the DRS instance type is set to primary/standby, duplicate or inconsistent data may occur in tables without primary keys when a primary/standby switchover is performed. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|---|
| Incremental synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. ● During migration of table-level objects, you are not advised to rename the tables. ● Replication of interval partition tables is not supported. ● The name of a primary key column cannot be changed. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. ● If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> ● After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> ● To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. ● The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-109 Synchronization task information

Table 5-111 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-110 Synchronization instance details

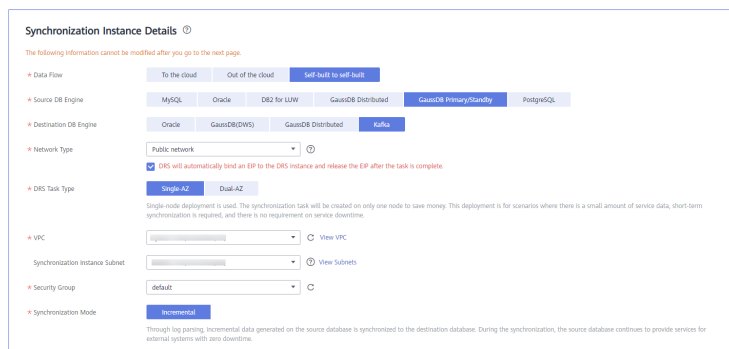


Table 5-112 Synchronization instance settings

| Parameter | Description |
|-----------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select GaussDB Primary/Standby . |
| Destination DB Engine | Select Kafka . |
| Network Type | Public network is used as an example. Available options: Public network and VPN or Direct Connect |
| DRS Task Type | Type of the DRS task. The value can be Single-AZ or Dual-AZ . <ul style="list-style-type: none"> Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p> |
| VPC | Select an available VPC. |

| Parameter | Description |
|---------------------------------|---|
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | <ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. <p>During synchronization, the source database continues to provide services for external systems with zero downtime.</p> |

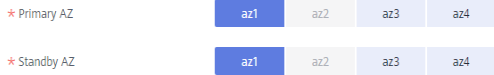
- Task Type

Figure 5-111 Task type



Table 5-113 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |

| Parameter | Description |
|-----------|---|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 5-112 AZ</p>  |

- Enterprise Project and Tags

Figure 5-113 Enterprise projects and tags

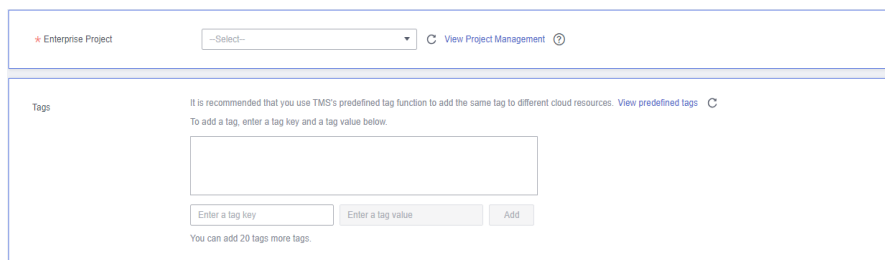


Table 5-114 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.


Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.


Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.


Figure 5-114 Source database information

Source Database

DR Synchronization 

IP Address or Domain Name 
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-115 Source database settings


| Parameter | Description |
|---------------------------|--|
| DR Synchronization | If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name . |
| IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-115 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 5-116 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. |

| Parameter | Description |
|-----------|--|
| 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 5-116 Synchronization mode

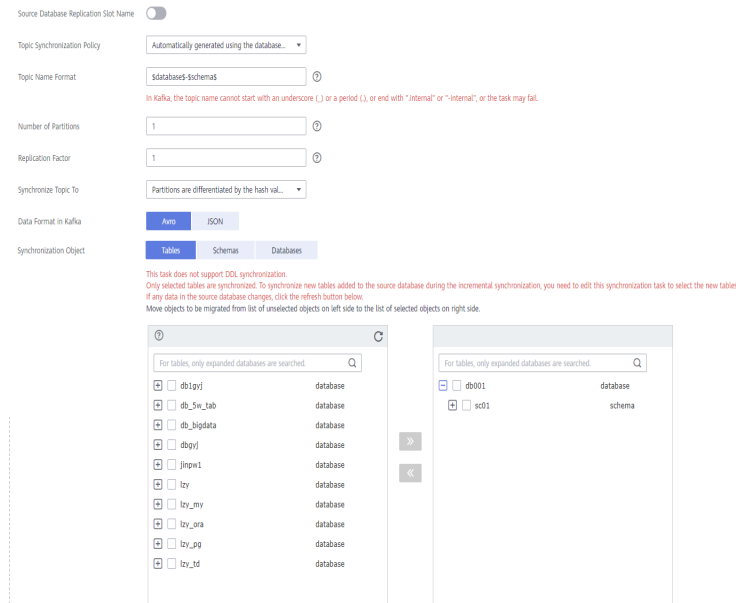



Table 5-117 Synchronization object

| 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 . |

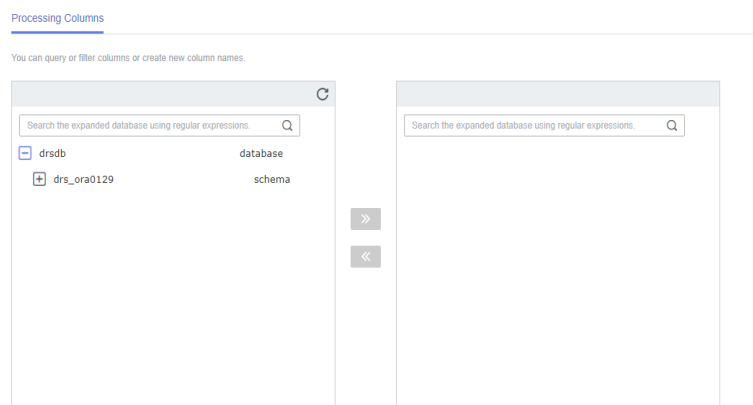
| Parameter | Description |
|----------------------|---|
| Topic Name Format | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Due to Kafka restrictions, a topic name can contain only ASCII characters, periods (.), underscores (_), and hyphens (-). If a topic name exceeds the limit, the topic fails to be created and the task is abnormal.</p> <p>If a topic name contains a database object name, ensure that the characters in the object name meet the Kafka topic naming requirements.</p> |
| Number of Partitions | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.</p> |
| Replication Factor | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p> |
| Synchronize Topic To | <p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> ● Partitions are differentiated by the hash values of database_name.schema_name.table_name: This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved. ● Partitions are identified by the hash values of the primary key: This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently. For a table without a primary key, if you select Partitions are identified by the hash values of the primary key, topics are synchronized to different partitions based on the hash values of database_name.schema_name.table_name. ● Partitions are differentiated by the hash values of database_name.schema_name: This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently. ● If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. |

| Parameter | Description |
|------------------------|---|
| Data Format in Kafka | <p>Select the format of data delivered to Kafka.</p> <ul style="list-style-type: none"> ● Avro: A binary encoded format that is efficient. ● JSON: JSON message format, which is easy to interpret but takes up more space. <p>For details, see Kafka Message Format.</p> |
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports table-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● To quickly select the desired database objects, you can use the search function. ● If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. ● If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Figure 5-117 Processing data



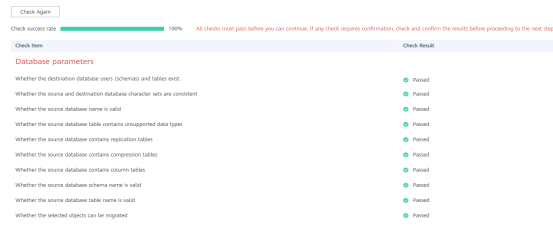
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-118 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-119 Task startup settings





Table 5-118 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |

| Parameter | Description |
|-----------------------------|--|
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.12 From GaussDB Primary/Standby to GaussDB Distributed

Supported Source and Destination Databases

Table 5-119 Supported databases

| Source DB | Destination DB |
|-------------------------|---|
| GaussDB primary/standby | GaussDB distributed NOTE The destination database version must be the same as or later than the source database version. |

Supported Synchronization Objects

Table 5-120 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-120 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Constraints |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-121](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-121 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|------------------|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 5-122 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Constraints |
|----------------------|--|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● 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. - If a unique index in the source database is an expression index, the destination distributed GaussDB may fail to create the index during full synchronization. As a result, the task fails. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |
| 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 | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 5-120 Synchronization task information

Table 5-123 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-121 Synchronization instance details

Table 5-124 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 GaussDB Distributed . |
| Network Type | Public network is used as an example. Available options: VPC, Public network, and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 5-122 Task type



Table 5-125 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-123 Enterprise projects and tags

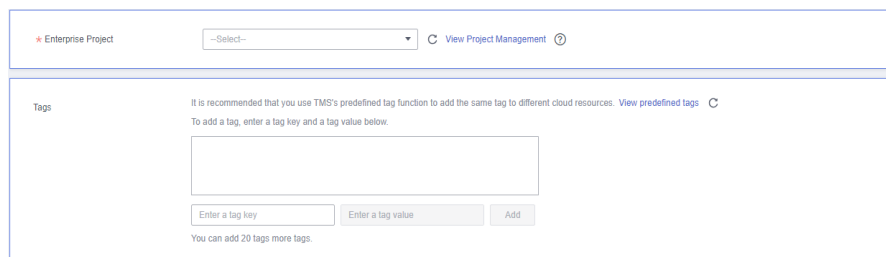


Table 5-126 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.


Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.


Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.


Figure 5-124 Source database information

Source Database

DR Synchronization 

IP Address or Domain Name 
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-127 Source database settings


| Parameter | Description |
|---------------------------|--|
| DR Synchronization | If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name . |
| IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-125 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance

Database Username


Database Password 

Table 5-128 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 |

| Parameter | Description |
|-------------------|--|
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-126 Synchronization mode

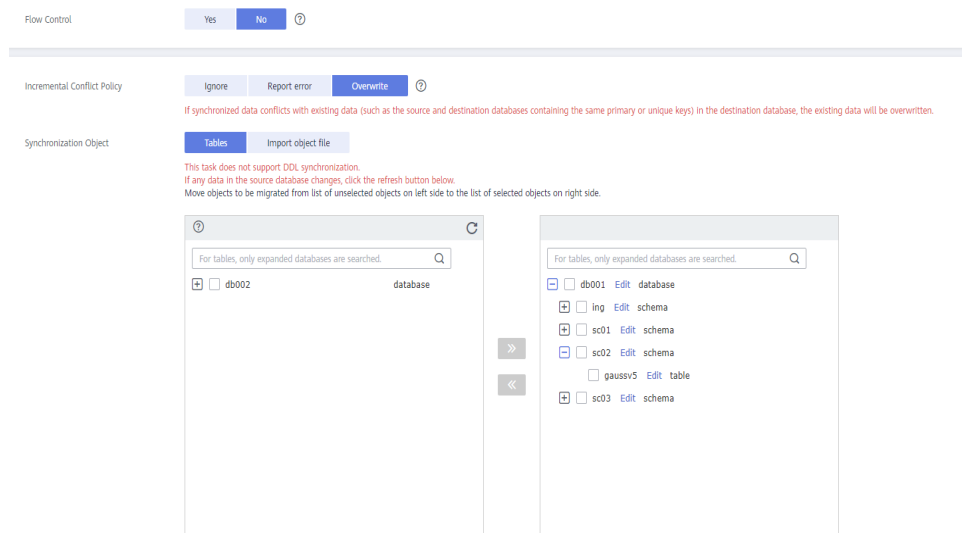
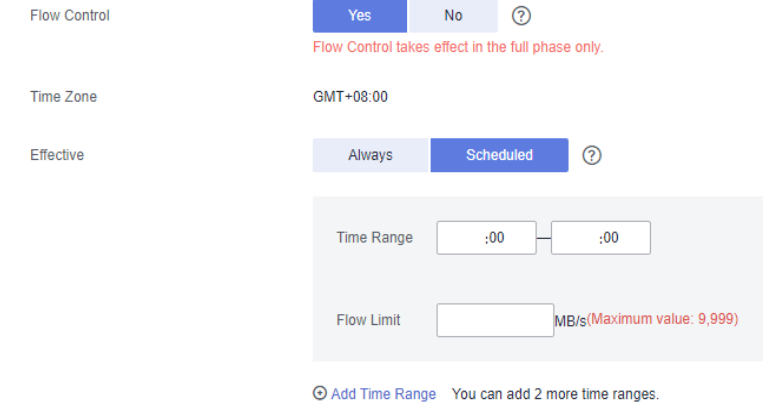



Table 5-129 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-127 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

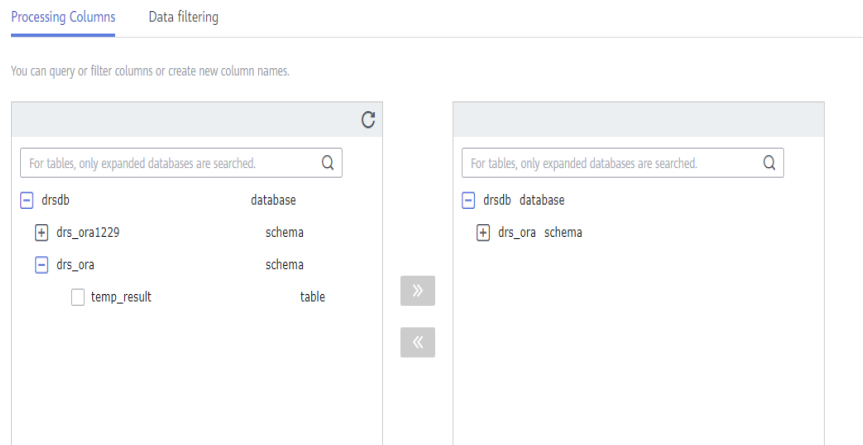
| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-128 Processing data



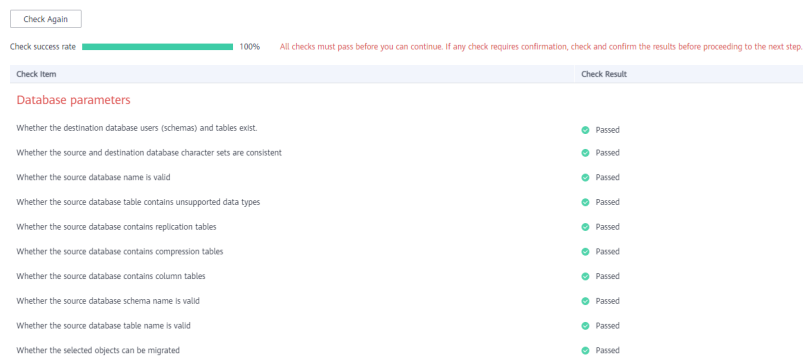
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-129 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-130 Task startup settings





Table 5-130 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.13 From GaussDB Primary/Standby to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-131 Supported databases

| Source DB | Destination DB |
|-------------------------|---|
| GaussDB primary/standby | <p>GaussDB primary/standby</p> <p>NOTE The destination database version must be the same as or later than the source database version.</p> |

Supported Synchronization Objects

Table 5-132 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-132 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Constraints |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-133](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-133 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|------------------|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 5-134 Precautions

| Type | Constraints |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' ?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Select full or full+incremental synchronization. If a table structure already exists in the destination database, ensure that the table structure is the same as that on the source database. If column processing is performed, ensure that the table structure is the same as that after column processing. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the |

| Type | Constraints |
|----------------------|---|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● 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. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key and unique key cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |
| 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 | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Primary/Standby. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 5-131 Synchronization task information

Table 5-135 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-132 Synchronization instance details

Table 5-136 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 GaussDB Primary/Standby . |
| Network Type | Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| 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. |

- Task Type

Figure 5-133 Task type



Table 5-137 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-134 Enterprise projects and tags




Table 5-138 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see Enterprise Management User Guide . To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i> . |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.


Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.


Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.


Figure 5-135 Source database information

Source Database

DR Synchronization 

IP Address or Domain Name 
Ensure that the entered addresses belong to the same DB instance.
Enter the IP addresses of all primary and standby instances.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-139 Source database settings

| Parameter | Description |
|---------------------------|--|
| DR Synchronization | If enabled, the source database is a cluster deployed in the DR mode. Information about all nodes in the DR cluster must be entered in IP Address or Domain Name . |
| IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance and enter the IP addresses of all primary and standby instances. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-136 Destination database information

Destination Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

Table 5-140 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). Example: 192.168.0.1:8000,192.168.0.2:8000 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-137 Synchronization mode

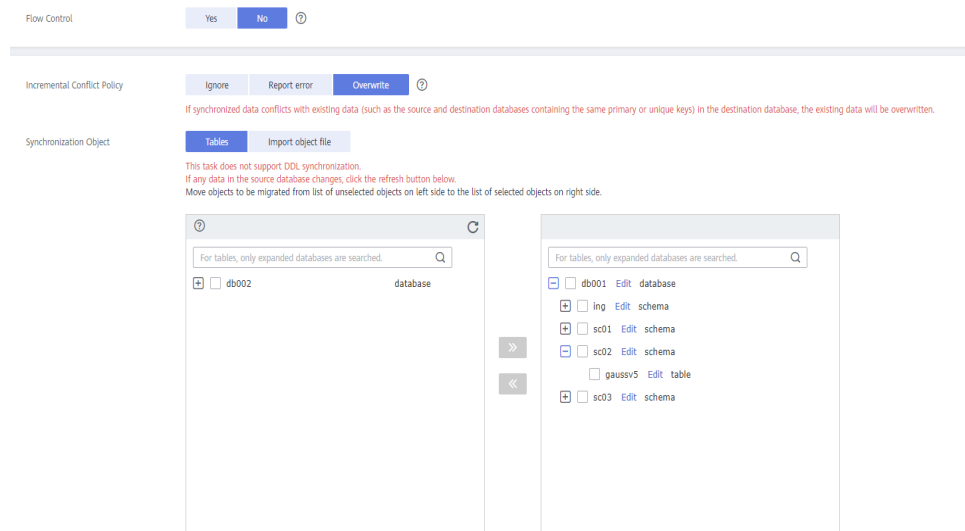
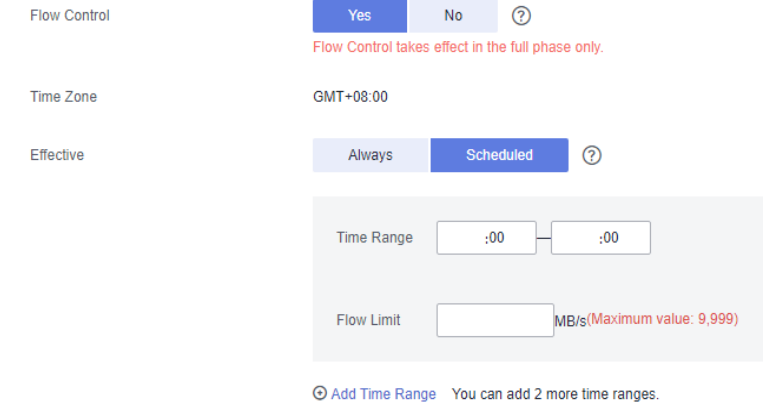



Table 5-141 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-138 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

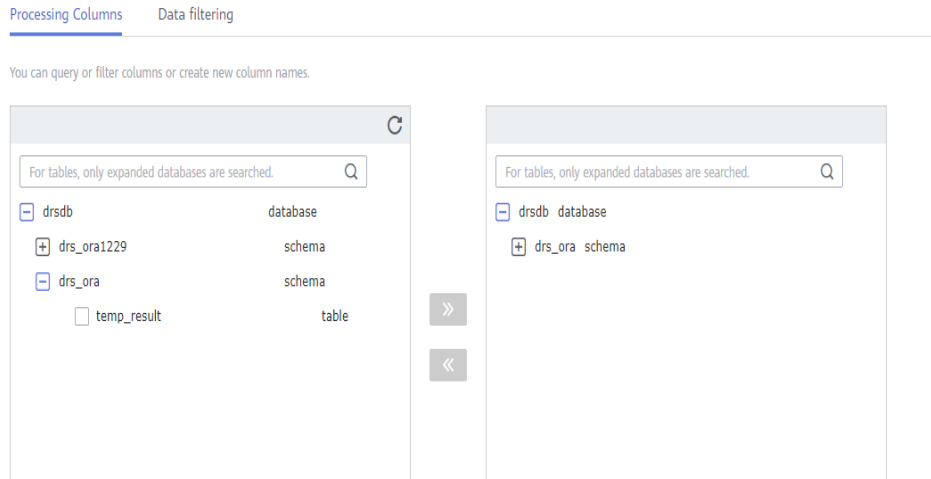
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-139 Processing data



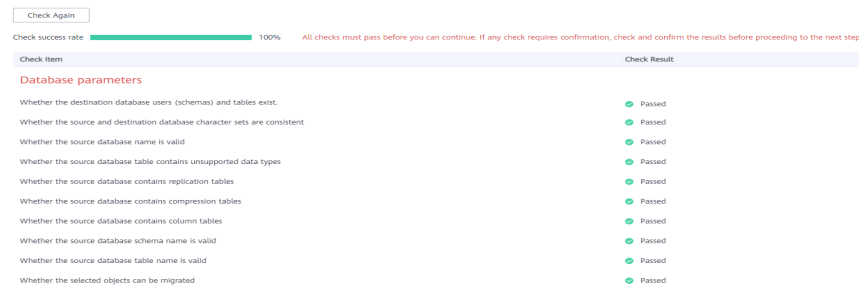
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-140 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-141 Task startup settings





Table 5-142 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|--|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.14 From GaussDB Distributed to MySQL

Supported Source and Destination Databases

Table 5-143 Supported databases

| Source DB | Destination DB |
|---------------------|---|
| GaussDB distributed | <ul style="list-style-type: none"> On-premises MySQL 5.5, 5.6, and 5.7 databases MySQL 5.5, 5.6, and 5.7 databases on an ECS MySQL 5.5, 5.6, and 5.7 databases on other clouds |

Supported Synchronization Objects

Table 5-144 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-144 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. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \\?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of

synchronization tasks require different permissions. For details, see [Table 5-145](#). 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 5-145 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|--|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user must have the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|--|--|
| | | details, see Connection and Port Description for Incremental Synchronization from GaussDB. | a GaussDB Database. <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | Required permissions: INSERT, DELETE, UPDATE, SELECT, and SHOW DATABASES | | |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 5-146 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \?!. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The source database supports the UTF8 character set, and the destination database supports the UTF8MB4 character set. If the source database uses the UTF8 character set and the destination database uses the UTF8MB3 character set, or if the source database uses the GBK 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 |

| Type | Restrictions |
|----------------------|--|
| | <p>requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● 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 names (case-sensitive) of objects, such as tables and schemas, in the source database must be the same as those in the destination database. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Before creating a DRS task, if concurrency control rules of SQL statements are configured for the destination database, the DRS task may fail. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| Full synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Do not execute any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. |

| 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 GaussDB Developer Guide. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Do not limit the synchronization speed during data comparison. ● MySQL and GaussDB process zero time (0000-00-00 00:00:00) in different ways. During value comparison, if there is zero time in the source or destination database, the comparison result shows that data is inconsistent. |

| Type | Restrictions |
|-----------------|--|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-142 Synchronization task information

Table 5-147 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-143 Synchronization instance details

Table 5-148 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 MySQL . |
| Network Type | Public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | <p>Available options: Full+Incremental, Incremental, and Full</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

| Parameter | Description |
|--------------------|---|
| Source DN Quantity | The value must be the same as the number of DNs in the distributed source database. |

- Task Type

Figure 5-144 Task type



Table 5-149 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-145 Enterprise projects and tags

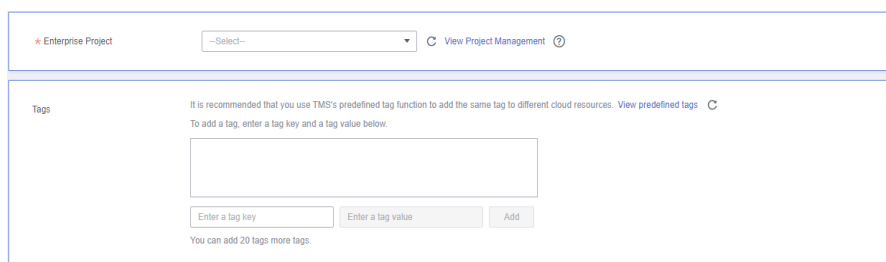


Table 5-150 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-146 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 5-151 Source database settings

| Parameter | Description |
|------------------------------|---|
| CN IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| DN IP Address or Domain Name | You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel. |

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-147 Destination database information

Destination Database

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

✔ Test successful

Table 5-152 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If SSL is disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-148 Synchronization mode

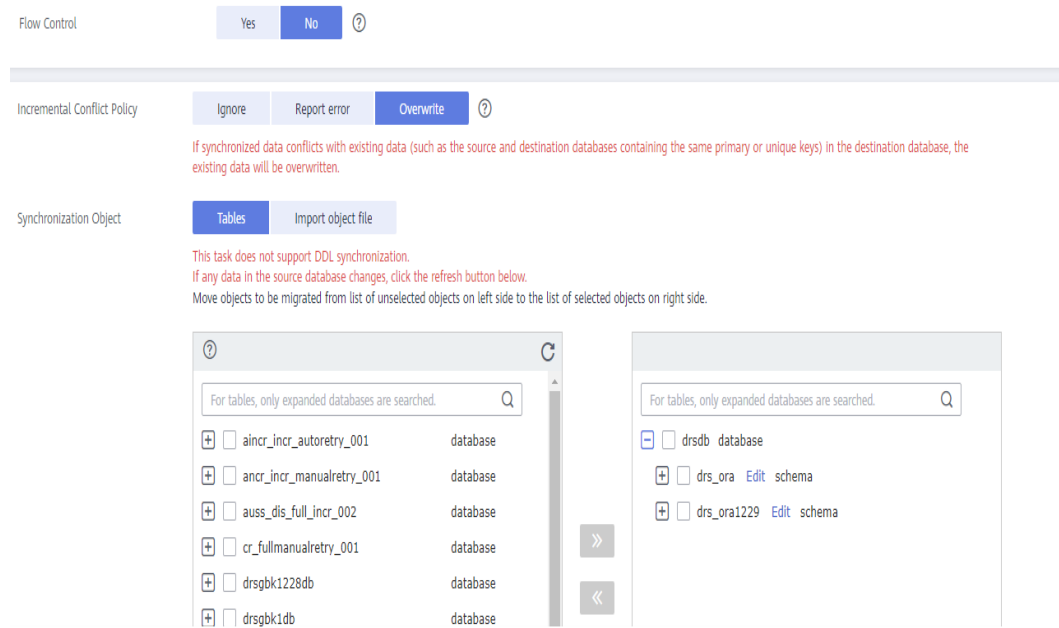
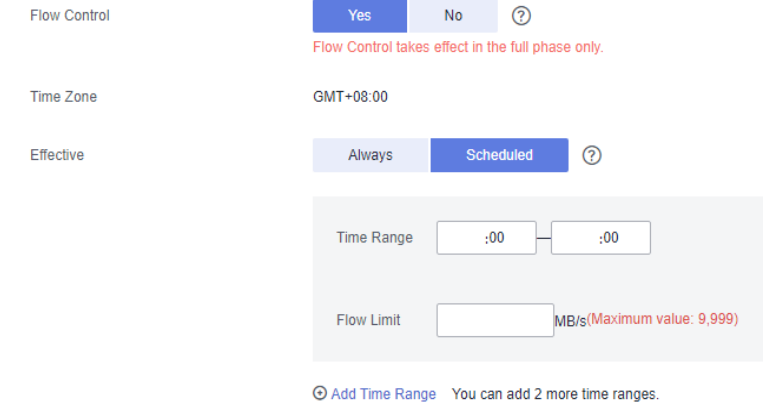



Table 5-153 Synchronization Object

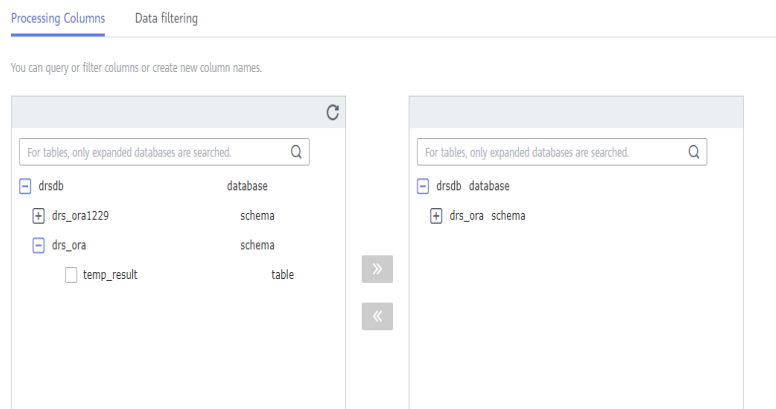
| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-149 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-150 Processing data



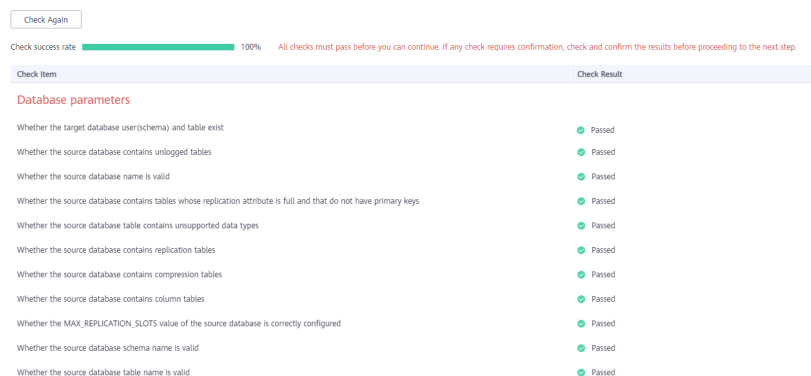
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-151 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-152 Task startup settings





Table 5-154 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.15 From GaussDB Distributed to Oracle

Supported Source and Destination Databases

Table 5-155 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 5-156](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-156 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. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' !?. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). - 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 5-157](#). 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 5-157 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|---|---|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | <p>a GaussDB Database.</p> <ul style="list-style-type: none"> Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |
| Destination database user | 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.

- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 5-158 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. - When the destination database is Oracle, dates ranging from 1582-10-05 to 1582-10-14 are not supported because data will be inconsistent between source and destination databases. - In the source GaussDB database, the date type for a BC leap year is 1, 5, or 9. In the destination Oracle database, the date type for a BC leap year is 0, 4, or 8. February 29 in a BC leap year in the source database is not supported in the destination database. As a result, the date fails to be parsed and the synchronization fails. - If a table to be synchronized has a composite unique constraint that can be null, data inconsistency or task |

| Type | Restrictions |
|----------------------|--|
| | <p>failure may occur due to the difference in the constraint range of null values in GaussDB and Oracle. For example, a table to be synchronized contains the UNIQUE(C1,C2) constraint that can be null. In GaussDB, there can be multiple data records whose c1='1' and c2=null at the same time. In Oracle, there can be only one data record whose c1='1' and c2=null. If one more data record is inserted again, a unique key conflict occurs.</p> <ul style="list-style-type: none"> - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |
| 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 GaussDB Developer Guide. ● For tables containing large fields (including blob, clob, nclob, long, and long raw), large fields are not used as where conditions for incremental data. You are advised to ensure that all fields except large fields have the unique constraint. ● If a table in the destination database contains both the lob and long types or both the lob and long raw types, and the length of the long or long raw type to be inserted or updated is greater than 4000, the insertion or update fails. You are advised to use the clob or blob type. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. |
| Synchronization comparison | <ul style="list-style-type: none"> ● You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● Data cannot be compared during full synchronization. ● Do not limit the synchronization speed during data comparison. |

| Type | Restrictions |
|-----------------|--|
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [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 5-153 Synchronization task information

Table 5-159 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-154 Synchronization instance details

Table 5-160 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 Oracle . |
| Network Type | Public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none"> - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

| Parameter | Description |
|--------------------|---|
| Source DN Quantity | The value must be the same as the number of DNs in the distributed source database. |

- Task Type

Figure 5-155 Task type



Table 5-161 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization . NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications . |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-156 Enterprise projects and tags

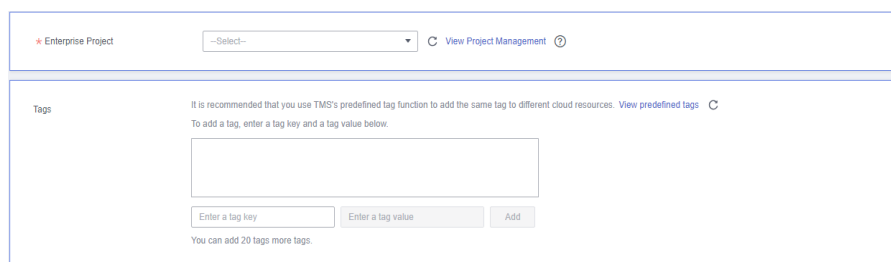


Table 5-162 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-157 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 5-163 Source database settings

| Parameter | Description |
|------------------------------|---|
| CN IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| DN IP Address or Domain Name | You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel. |

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-158 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 5-164 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | The IP address or domain name of the destination database. NOTE For a RAC cluster, use a scan IP address and specify Service Name to improve access performance. |
| Port | The port of the destination database. Range: 1 - 65535 |
| Database Service Name | Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |
| SSL Connection | SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none"> • The maximum size of a single certificate file that can be uploaded is 500 KB. • If SSL is disabled, your data may be at risk. |

 NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the objects to be synchronized, and then click **Next**.

Figure 5-159 Synchronization mode

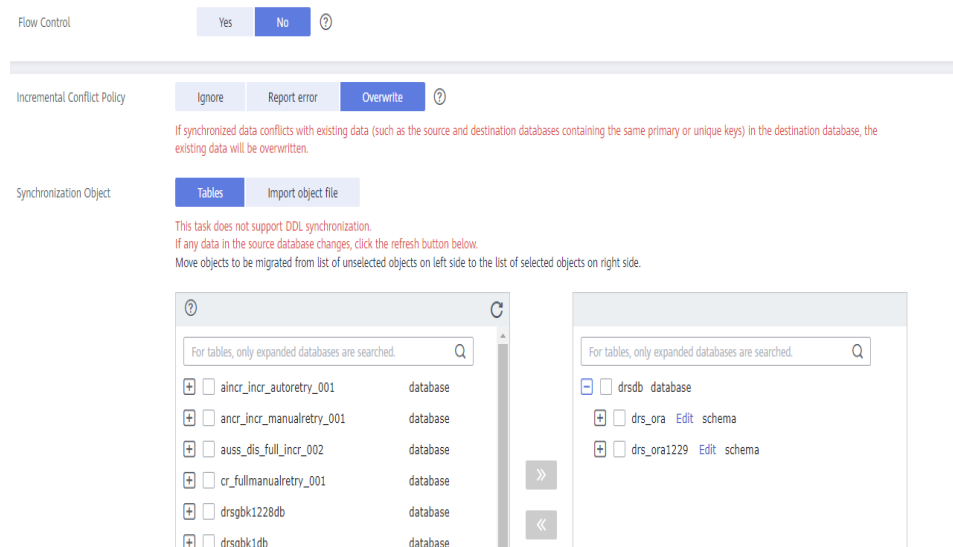
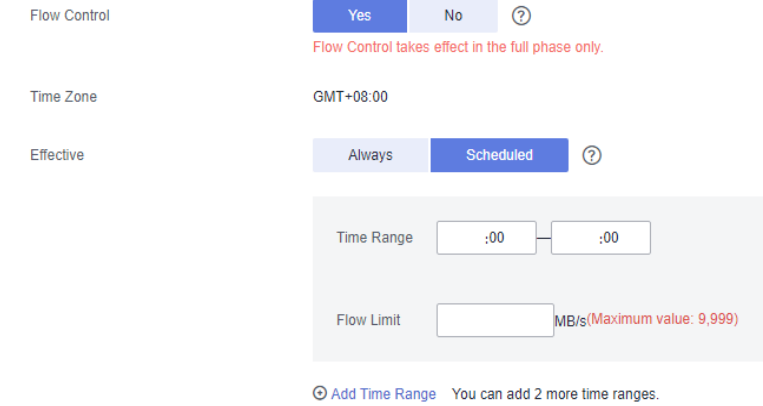



Table 5-165 Synchronization Object

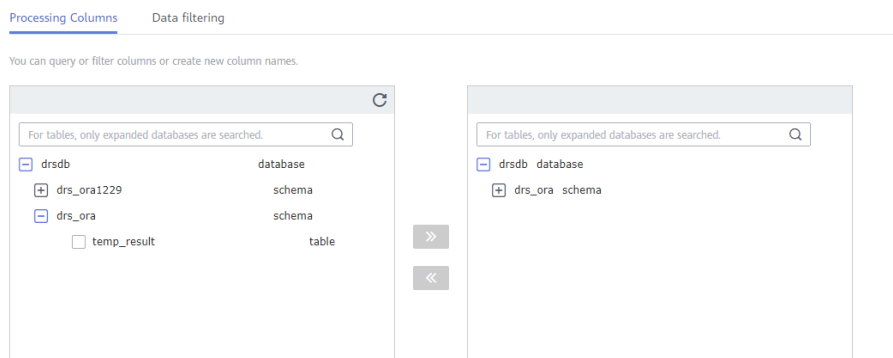
| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-160 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Report error The synchronization task will be stopped and fail. ● Overwrite Conflicting data will be overwritten. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. ● The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-161 Processing data



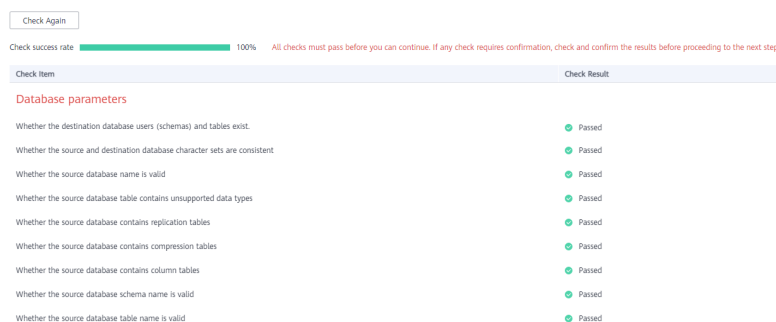
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-162 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-163 Task startup settings

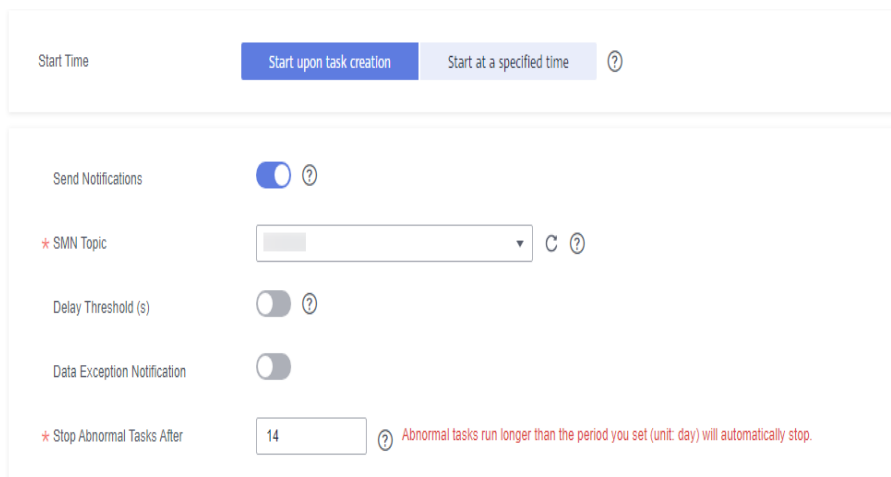



Table 5-166 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.16 From GaussDB Distributed to Kafka

Supported Source and Destination Databases

Table 5-167 Supported databases

| Source DB | Destination DB |
|---------------------|---------------------|
| GaussDB distributed | Kafka 0.11 or later |

Supported Synchronization Objects

Table 5-168 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-168 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 • The supported fields are BIGINT, BIT, BIT VARYING, BLOB, BOOLEAN, BYTEA, CHARACTER, CHARACTER VARYING, CLOB, DATE, DOUBLE PRECISION, INTEGER, MONEY, NUMBER, NUMERIC, NVARCHAR2, RAW, REAL, SMALLDATETIME, SMALLINT, TEXT, TIME WITH TIME ZONE, TIME WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITHOUT TIME ZONE and TINYINT. • Table-level synchronization, schema-level synchronization, and database-level synchronization are supported. <ul style="list-style-type: none"> - Only DML statements of the selected table can be synchronized. - Databases without schemas cannot be synchronized. - Schemas without tables cannot be synchronized. - Column-store tables, compressed tables, delay tables, and temporary tables cannot be synchronized. Do not synchronize unlogged tables. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' \\?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-169](#).

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 5-169 Database user permission

| Type | Incremental |
|----------------------|---|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. • The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. • Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [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 5-170 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Set the REPLICA IDENTITY attribute of a table without a primary key to FULL, or add a primary key to the table. - Set the REPLICA IDENTITY attribute of the table that has a primary key to FULL. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\` \?! ● Destination database requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - If a logical replication slot fails to be created or does not exist due to a long transaction, you can reset the task and then restart it. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|--|
| Incremental synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • During migration of table-level objects, you are not advised to rename the tables. • Replication of interval partition tables is not supported. • The name of a primary key column cannot be changed. • After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. • Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding of GaussDB Developer Guide." • If an incremental synchronization task is suspended or resumed due to an exception, there may be duplicate data in the destination Kafka. Use the id field in the Kafka data for data deduplication. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
 - Task information description

Figure 5-164 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, there are four input fields:

- Region:** A dropdown menu with a globe icon on the left and a downward arrow on the right.
- Project:** A dropdown menu with a downward arrow on the right.
- Task Name:** A text input field containing "DRS-6131" and a question mark icon on the right.
- Description:** A text area with a question mark icon on the right and a character count "0/256" at the bottom right.

 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."

Table 5-171 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-165 Synchronization instance details

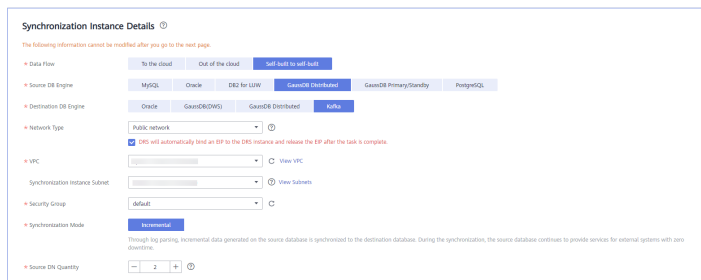


Table 5-172 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Choose Self-built to self-built . |
| 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 | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> <p>If you select VPN or Direct Connect for Network Type, you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created.</p> |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <ul style="list-style-type: none"> Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. During synchronization, the source database continues to provide services for external systems with zero downtime. |
| Source DN Quantity | The value must be the same as the number of DNs in the distributed source database. |

- Task Type

Figure 5-166 Task type



Table 5-173 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-167 Enterprise projects and tags

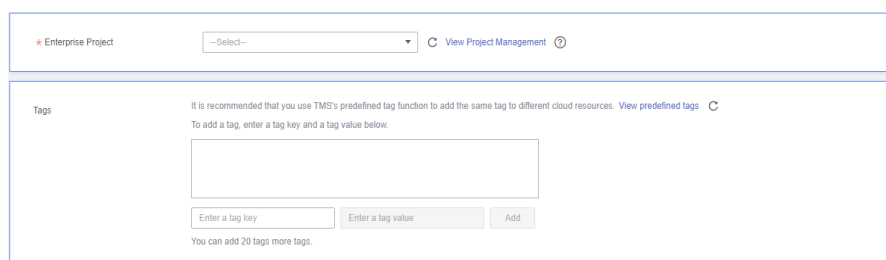


Table 5-174 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-168 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 5-175 Source database settings


| Parameter | Description |
|------------------------------|---|
| CN IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| DN IP Address or Domain Name | You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel. |

NOTE

The username and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.

Figure 5-169 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 5-176 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080. |
| Connection Method | Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication . |

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Figure 5-170 Synchronization mode

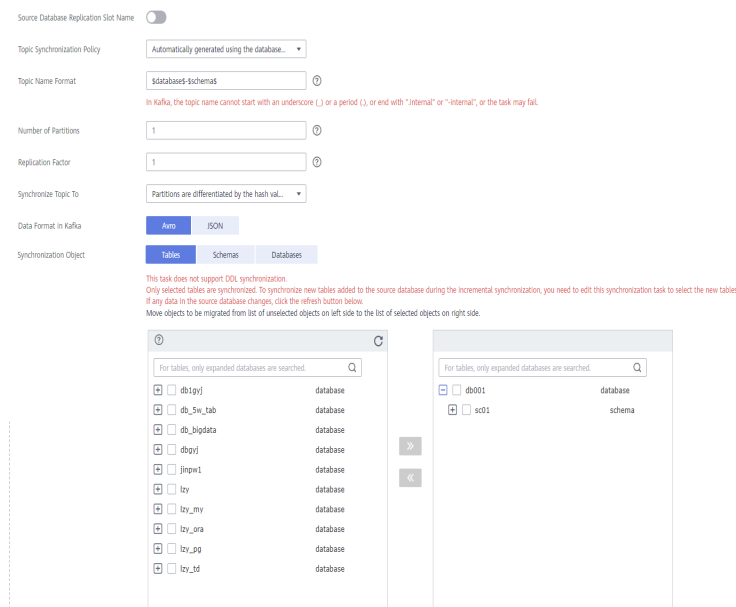



Table 5-177 Synchronization Object

| 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. |

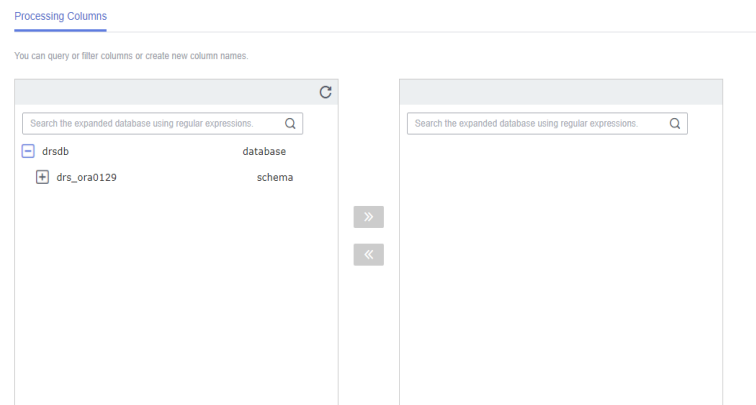
| Parameter | Description |
|----------------------|---|
| Number of Partitions | <p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput 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-, schema-, and database-level synchronization. You can select data for synchronization based on your service requirements.</p> <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set the corresponding rules by referring to [Processing Data](#).

Figure 5-171 Processing data



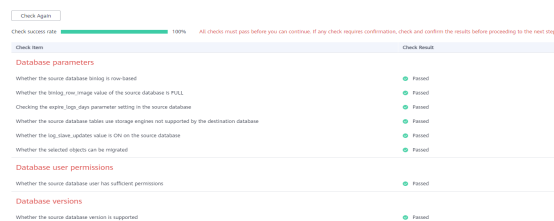
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-172 Pre-check



 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-173 Task startup settings

Table 5-178 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.17 From GaussDB Distributed to GaussDB Distributed

Supported Source and Destination Databases

Table 5-179 Supported databases

| Source DB | Destination DB |
|---------------------|---|
| GaussDB distributed | GaussDB distributed NOTE The destination database version must be the same as or later than the source database version. |

Supported Synchronization Objects

Table 5-180 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-180 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?! The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Constraints |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-181](#). 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 5-181 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|--|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. To compare content, the MONADMIN permission is required. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|--|--|------------------|
| | | enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB . | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • To use value comparison, you must have the MONADMIN permission. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

 CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
- To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.

-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
- To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 5-182 Precautions

| Type | Constraints |
|-----------------|--|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: If incremental synchronization or full+incremental synchronization is selected: <ul style="list-style-type: none"> - The wal_level parameter of the source database is set to logical. - The enable_slot_log parameter of the source database is set to on. - The max_replication_slots value of the source database must be greater than the number of used replication slots. - Add a primary key to the table that does not have a primary key, or set REPLICATION IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"/?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records |

| Type | Constraints |
|------|---|
| | <p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - 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. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|--|
| 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. |
| 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 GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |

| Type | Constraints |
|----------------------------|--|
| Synchronization comparison | <ul style="list-style-type: none"> • You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • The topology structure of the source database must be the same as that of the destination database. Otherwise, the content comparison function is unavailable. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • Data cannot be compared during full synchronization. • Do not limit the synchronization speed during data comparison. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> • After the task is complete, the streaming replication slot created in the source database is automatically deleted. • If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> • To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. • To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. • The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).
- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-174 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-8131 [Text input]

Description: [Text area]

0/256

Table 5-183 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-175 Synchronization instance details

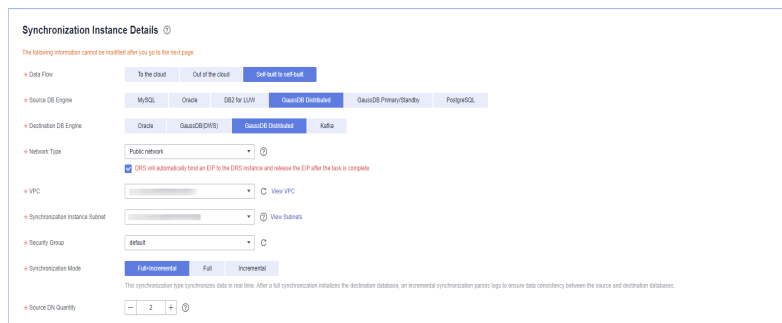


Table 5-184 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select GaussDB Distributed . |
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | Public network is used as an example. Available options: VPC, Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. If you select VPN or Direct Connect for Network Type , you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |
| Source DN Quantity | The value must be the same as the number of DNs in the distributed source database. |

- Task Type

Figure 5-176 Task type



Table 5-185 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-177 Enterprise projects and tags




Table 5-186 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-178 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 5-187 Source database settings

| Parameter | Description |
|------------------------------|---|
| CN IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| DN IP Address or Domain Name | You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel. |

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-179 Destination database information

Destination Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password ⓘ

Table 5-188 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-180 Synchronization mode

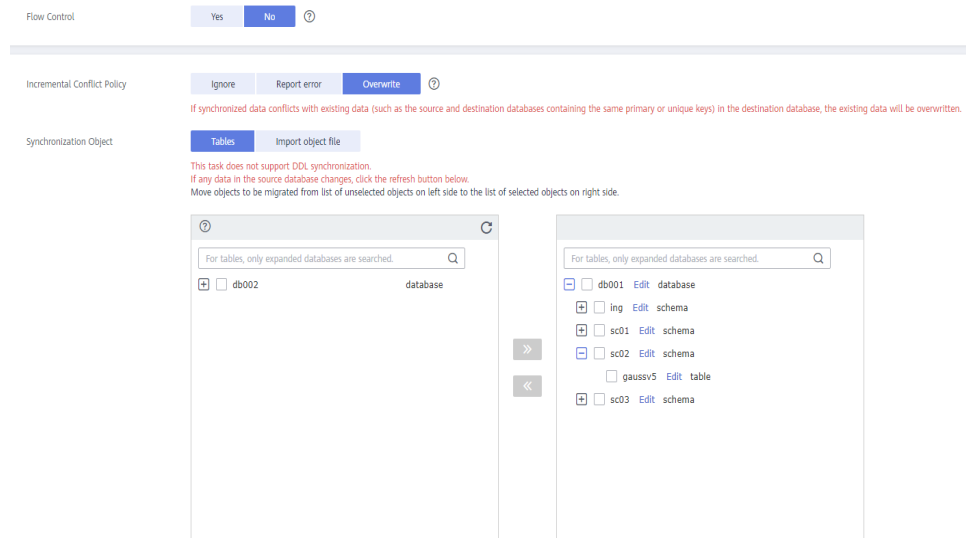
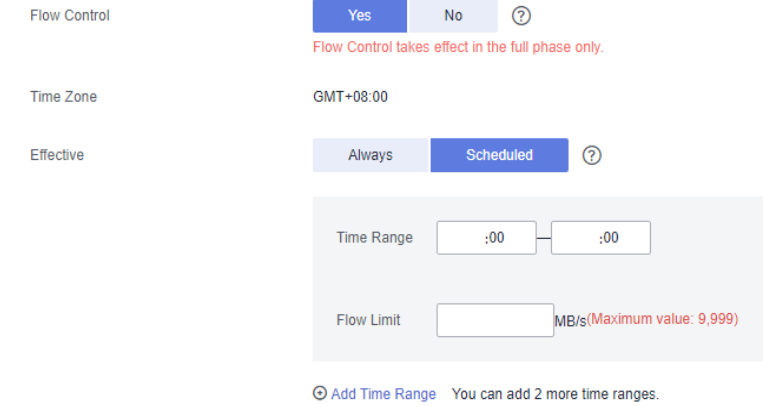



Table 5-189 Synchronization Object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-181 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |


| Parameter | Description |
|-----------------------------|---|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-182 Processing data



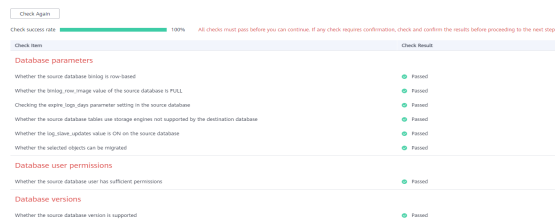
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-183 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-184 Task startup settings





Table 5-190 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.18 From GaussDB Distributed to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-191 Supported databases

| Source DB | Destination DB |
|---------------------|---|
| GaussDB distributed | <p>GaussDB primary/standby</p> <p>NOTE</p> <ul style="list-style-type: none"> The destination database version must be the same as or later than the source database version. |

Supported Synchronization Objects

[Table 5-192](#) lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-192 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, temporary tables, level-2 partitioned tables, and replication tables without primary keys cannot be synchronized. Do not synchronize unlogged tables in the incremental phase. - The default values of generated columns cannot be synchronized. - A user-defined table that has no primary key and is a partitioned table cannot be synchronized. Otherwise, the task may fail. - System schemas (pg_toast, cstore, snapshot, sys, dbms_job, dbms_perf, pg_catalog, information_schema, utl_file, dbms_output, dbms_random, utl_raw, dbms_sql, dbms_lob, dbe_perf, pkg_service, pkg_util, dbe_file, dbe_random, dbe_output, dbe_raw, dbe_sql, dbe_lob, dbe_task, blockchain, db4ai, dbe_pldebugger, sqladvisor, dbe_application_info, dbe_match, dbe_pldeveloper, dbe_scheduler, dbe_session, dbe_utility, dbe_sql_util, dbe_xml, dbe_xmldom, dbe_xmlparser, dbe_compression, dbe_heat_map, dbe_ilm, dbe_ilm_admin, prvt_ilm, dbe_profiler, dbe_stats, rdsBackup, rdsMetric, rdsRepl and public) cannot be synchronized. - System tables (redis_progress_detail, redis_status, pgxc_redsitb and redis_progress in the public schema) cannot be synchronized. - The database name, schema name, and table name cannot contain special characters /<.>\\' ?!. The column name cannot contain double quotation marks ("), single quotation marks ('), or periods (.). |

| Type | Constraints |
|------|---|
| | <ul style="list-style-type: none">- If you select tables by importing an object file, ensure that the imported table exists in the source database or is visible to the synchronization user. |

Database User Permission Requirements

Before you start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-193](#). 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 5-193 Database user permission

| Type | Full | Incremental | Full+Incremental |
|----------------------|---|---|--|
| Source database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The CONNECT permission for databases, USAGE permission for schemas, SELECT or UPDATE permission for tables, the UPDATE permission for locking tables without primary keys, and SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, and the SELECT permission for tables. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> The REPLICATION permission or the permission inherited from the built-in role gs_role_replication, the CONNECT permission for databases, the USAGE permission for schemas, the SELECT or UPDATE permission for tables, the UPDATE permission for locking tables that do not have primary keys, and the SELECT permission for sequences. The user must have the remote connection permission. For details about how to assign the remote connection permission to a database user, see Configuring Remote Connection to a GaussDB Database. Ensure that the connection ports are enabled in the security group and firewall policies. For details, see Connection and Port Description for Incremental Synchronization from GaussDB. |

| Type | Full | Incremental | Full+Incremental |
|---------------------------|---|---|------------------|
| | | Connection and Port Description for Incremental Synchronization from GaussDB. | |
| Destination database user | <p>The user has the sysadmin role or the following minimum permissions:</p> <ul style="list-style-type: none"> • The CONNECT and CREATE permissions for databases, the USAGE and CREATE permissions for schemas, and the INSERT, UPDATE, DELETE, SELECT, ALTER, INDEX, and COMMENT permissions for tables. • If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. | | |

Suggestions

 **CAUTION**

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.

- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- For more information about the impact of DRS on databases, see [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 5-194 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 REPLICIA IDENTITY to FULL for the table that does not have a primary key. ● Source database object requirements: <ul style="list-style-type: none"> - The source database must be a GaussDB distributed instance. - The names of the source database, schema, and table to be synchronized cannot contain special characters /<.>\\' \"/?! ● Destination database parameter requirements: <ul style="list-style-type: none"> - The character set of the destination database must be the same as that of the source database. - Ensure that the source and destination databases are compatible with each other. ● Destination database object requirements: <ul style="list-style-type: none"> - Before the synchronization, ensure that the corresponding database has been created in the destination instance. - If you select incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - If you select full or full+incremental synchronization, ensure that the table structure of the destination database is the same as that of the source database. If columns are processed in the DRS instance, ensure that the table structure in the destination database is the same as the processed table structure. - Triggers enabled in the destination database cannot be associated with synchronization tables. - The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records |

| Type | Constraints |
|------|---|
| | <p>are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.</p> <p>Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.</p> <ul style="list-style-type: none"> ● Other notes: <ul style="list-style-type: none"> - SSL is not supported. - 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. - Restricted by the logical replication capability of GaussDB, up to 20 logical replication slots can be enabled for decoding on the same instance. |

| Type | Constraints |
|-----------------------------|---|
| 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. |
| Incremental synchronization | <ul style="list-style-type: none"> ● Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. ● Before a task enters the incremental synchronization phase, ensure that long-running transactions are not started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. ● Do not execute any DDL statement in the source database. Restricted by the logical replication function of GaussDB, DDL statements cannot be synchronized. Otherwise, data may be inconsistent or the task may fail. ● Do not change the REPLICA IDENTITY value of a table in the source database. Otherwise, incremental data may be inconsistent or the task may fail. ● Do not write data to the destination database. Otherwise, data may be inconsistent. ● Replication of interval partition tables is not supported. ● After a DDL statement is executed in a transaction, the DDL statement and subsequent statements are not synchronized. ● Logical log decoding is restricted by the decoding capability of the GaussDB kernel. For details about the restrictions, see the precautions in "Logical Decoding" of GaussDB Developer Guide. |
| Data processing | <ul style="list-style-type: none"> ● During column processing, the primary key, unique key, and distribution column cannot be filtered out. ● If table columns are renamed or filtered, the conditional expressions of the partial index in the index definition are ignored, expression columns in common indexes are excluded, and unique indexes that contain expression columns are excluded. ● If table columns are filtered, the corresponding columns in the common index and foreign key are also filtered. |

| Type | Constraints |
|----------------------------|--|
| Synchronization comparison | <ul style="list-style-type: none"> You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. Data cannot be compared during full synchronization. Do not limit the synchronization speed during data comparison. |
| Stopping a task | <p>Stop a task normally.</p> <ul style="list-style-type: none"> After the task is complete, the streaming replication slot created in the source database is automatically deleted. If the synchronization table in the destination database has been associated with a sequence when the task stops, DRS automatically resets the sequence value of the destination database. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. <p>Forcibly stop a task.</p> <ul style="list-style-type: none"> To forcibly stop a synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of GaussDB Distributed. To forcibly stop a synchronization task, delete the temporary tables without primary keys whose names are prefixed with drs_ in the destination database. The naming rule of a replication slot depends on the database kernel version (you can run the select working_version_num(); command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the task node id is *** log on the Synchronization Logs page. |

Prerequisites

- You have logged in to the DRS console.
- Your account balance is greater than or equal to \$0 USD.
- For details about the DB types and versions supported by real-time synchronization, see [Supported Databases](#).

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see [Agency Management](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Figure 5-185 Synchronization task information

Table 5-195 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| 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 5-186 Synchronization instance details

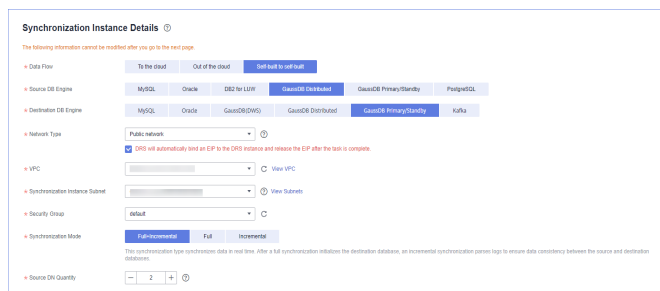


Table 5-196 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select GaussDB Distributed . |
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | Public network is used as an example. Available options: VPC , Public network , and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. If you select VPN or Direct Connect for Network Type , you can enter a specified IP address. When creating multiple tasks at the same time, do not specify the same group of unused IP addresses. Otherwise, the tasks fail to be created. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>Available options: Full+Incremental, Full, and Incremental. Full+Incremental is used as an example.</p> <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full In this mode, data is synchronized from the source to the destination at a time. - Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |
| Source DN Quantity | The value must be the same as the number of DNs in the distributed source database. |

- Task Type

Figure 5-187 Task type



Table 5-197 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-188 Enterprise projects and tags

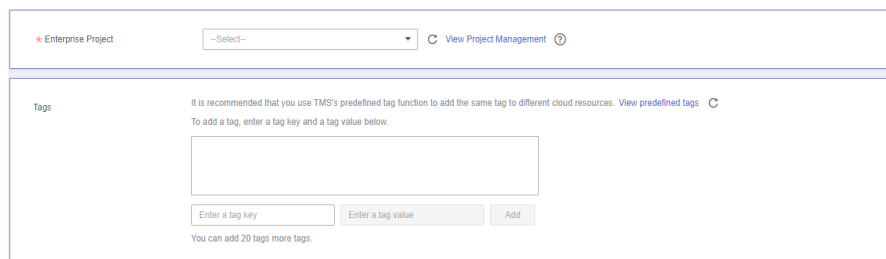


Table 5-198 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public

network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-189 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 5-199 Source database settings

| Parameter | Description |
|------------------------------|---|
| CN IP Address or Domain Name | Enter the IP addresses of GaussDB. Ensure that the entered IP addresses or domain names belong to the same instance. |
| Database Username | The username for accessing the source database. |
| Database Password | The password for the database username. |
| DN IP Address or Domain Name | You can log in to the CN node and run the SQL statement to query the DN IP address. The resource types of different database instances are different. Therefore, you are advised to contact O&M personnel. |

NOTE

The username and password of the source database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Figure 5-190 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



Test successful

Table 5-200 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). For example: 192.168.0.1:8080,192.168.0.2:8080. |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

NOTE

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization objects, and click **Next**.

Figure 5-191 Synchronization mode

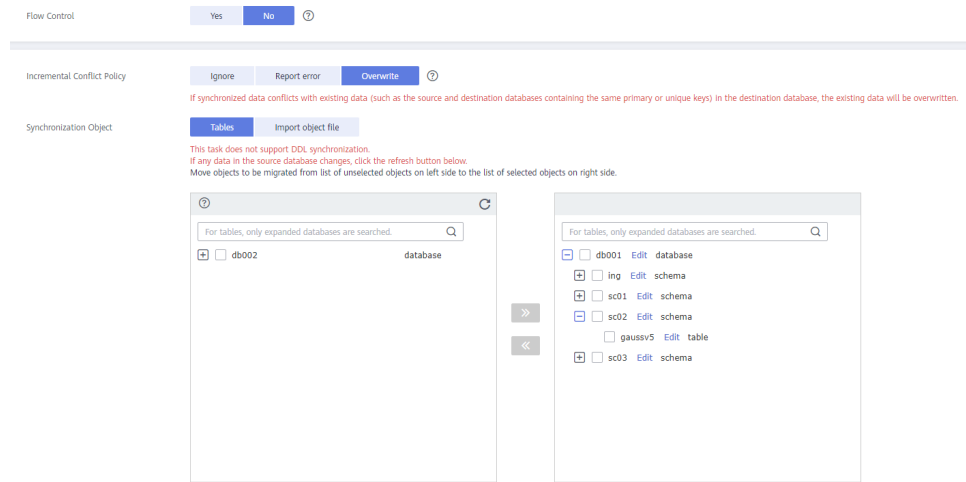
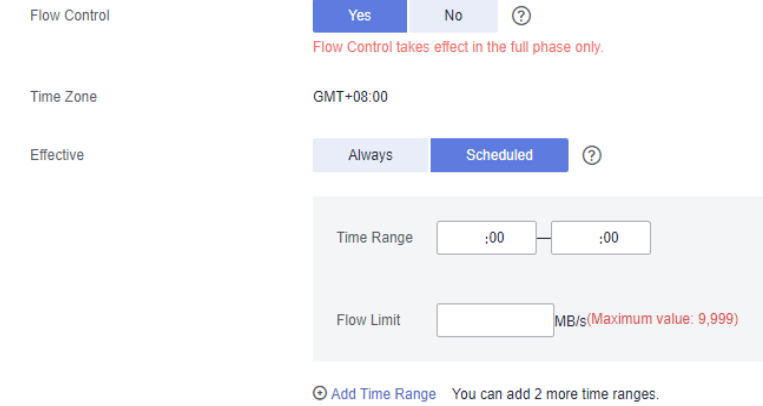



Table 5-201 Synchronization object

| Parameter | Description |
|--------------|--|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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 5-192 Flow control</p>  <ul style="list-style-type: none"> No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none"> The flow control mode takes effect only in the full synchronization phase. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode. |

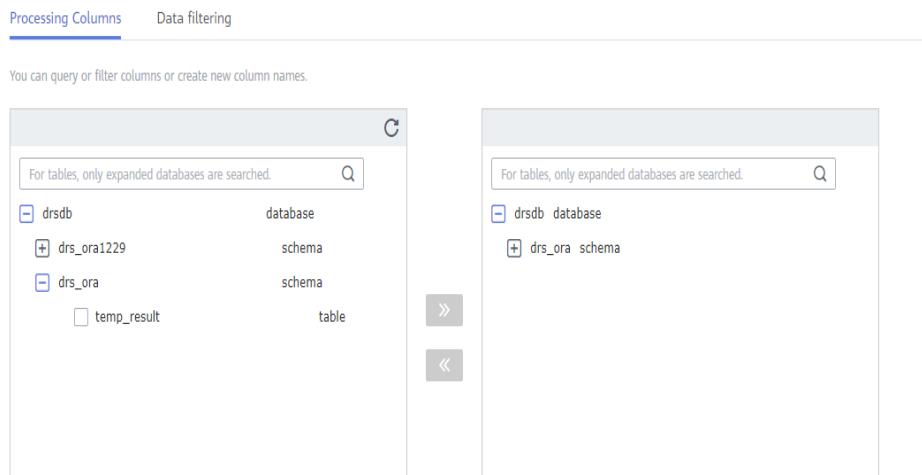
| Parameter | Description |
|-----------------------------|--|
| Incremental Conflict Policy | <p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> ● Ignore The system will skip the conflicting data and continue the subsequent synchronization process. ● Overwrite Conflicting data will be overwritten. ● Report error The synchronization task will be stopped and fail. <p>If the synchronized data conflicts with the existing data in the destination database, selecting Ignore cannot ensure data consistency. To ensure data consistency, select Overwrite.</p> |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. - The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.

- If you need to process columns or filter out data, set the corresponding rules by referring to [Processing Data](#).

Figure 5-193 Processing data



Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-194 Pre-check

Check success rate 100% All checks must pass before you can continue. If any check requires confirmation, check and confirm the results before proceeding to the next step.

| Check Item | Check Result |
|---|---|
| Database parameters | |
| Whether the source and destination databases are compatible with each other | ● Confirm Confirm Details |
| Whether the source database table contains unsupported field types | ● Passed |
| Source table replication attribute check | ● Passed |
| Whether the source and destination database character sets are consistent | ● Passed |
| Whether the destination database contains the configured databases | ● Passed |
| Whether the target database users(schema) and table exist | ● Passed |
| Whether the selected table contains delay constraints | ● Passed |
| Whether the source database contains unlogged tables | ● Passed |

NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-195 Task startup settings





Table 5-202 Task startup settings

| Parameter | Description |
|---------------------|--|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |

| Parameter | Description |
|-----------------------------|---|
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.19 From DB2 for LUW to GaussDB Primary/Standby

Supported Source and Destination Databases

Table 5-203 Supported databases

| Source DB | Destination DB |
|---|-------------------------|
| DB2 for LUW 9.7, 10.1, 10.5, 11.1, and 11.5 | GaussDB primary/standby |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - For more information about the impact of DRS on databases, see [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 and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-204 Precautions

| Type | Restrictions |
|----------------------|--|
| Database permissions | <ul style="list-style-type: none"> • Source database <ul style="list-style-type: none"> - To perform full synchronization, the CONNECT and DATAACCESS permissions are required. - To perform full+incremental synchronization, the DBADM permission is required. - If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the Datacapture attribute of the table to be synchronized is Y. • Destination database <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. - Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. - Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> • During full synchronization, tables, common indexes, primary 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 maximum precision supported by the timestamp type is 6. • During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. • The LOB type supported by incremental synchronization cannot exceed 10 MB. • Tables whose default values contain expression functions cannot be synchronized. • Temporary tables in the source database cannot be synchronized. • DDL statements executed in the source database cannot be synchronized. • For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. • If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency. |
| Source database | <ul style="list-style-type: none"> • Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or special characters .>`<`\, ?!". Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. • Only the GBK and UTF8 character sets can be synchronized. |

| Type | Restrictions |
|----------------------|--|
| Destination database | <ul style="list-style-type: none">• The destination database must be a GaussDB primary/standby instance.• Ensure that a database named in lowercase letters has been created in destination database.• Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.• The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none">• Do not use an imprecise value type as the primary key in the database because the synchronization of UPDATE and DELETE statements in the DRS incremental scenario will be affected.• ARM VMs are not supported.• If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization.• If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database.• After a task is started, the detach operation on a partition table may cause data inconsistency.• During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• During the synchronization, the source database cannot contain uncommitted DDL transactions.• Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database.• When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table.• During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database.• Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent.• During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated.• You can add additional objects in the incremental synchronization phase. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-196 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." The Task Name field has a character count of 0/256, and the Description field has a character count of 0/256.

Table 5-205 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-197 Synchronization instance details

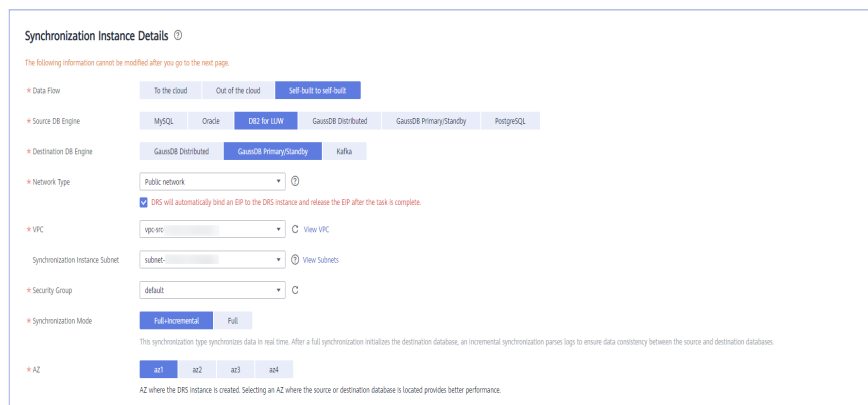


Table 5-206 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select DB2 for LUW. |
| Destination DB Engine | Select GaussDB Primary/Standby . |
| Network Type | The public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. |
| 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 Project and Tags

Figure 5-198 Enterprise projects and tags




Table 5-207 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-199 Source database information


Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password 

Database Name

SSL Connection

Synchronize Driver Synchronized.jcc-11.5.6.0(2).jar

Table 5-208 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. |


| Parameter | Description |
|--------------------|--|
| Database Password | The password for the source database username. |
| Database Name | The name of the database to which the synchronization objects belong in the source DB2 for LUW. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If the SSL certificate is not used, your data may be at risk. Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures. |
| Synchronize Driver | You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection. |

 **NOTE**

The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.


Figure 5-200 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-209 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). Example: 192.168.0.1:8000,192.168.0.2:8000 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 5-201 Synchronization mode

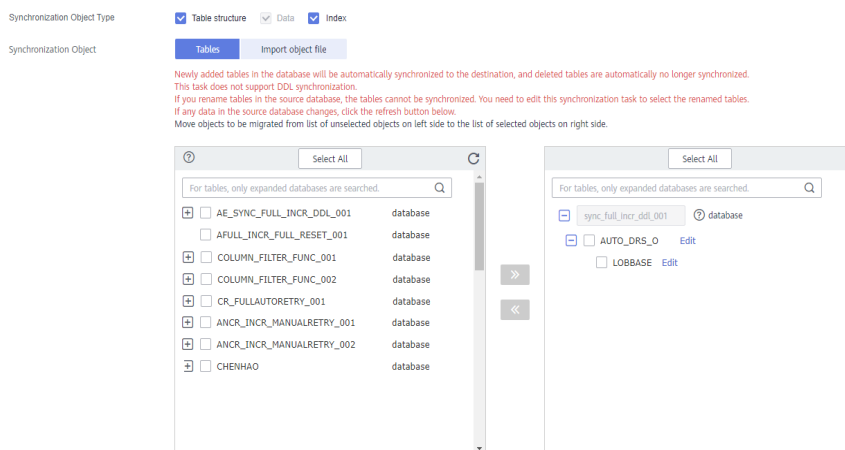



Table 5-210 Synchronization Object

| Parameter | Description |
|-----------------------------|--|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

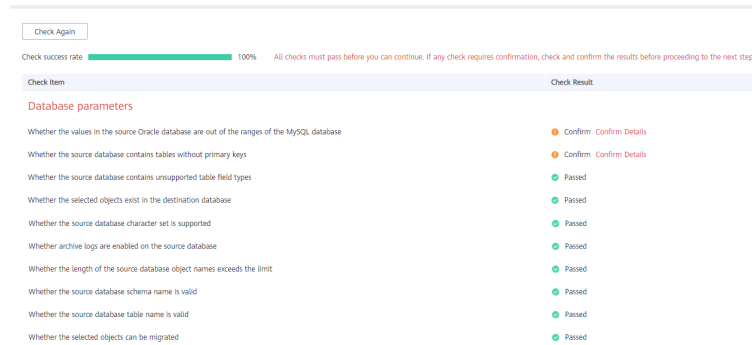
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-202 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-203 Task startup settings

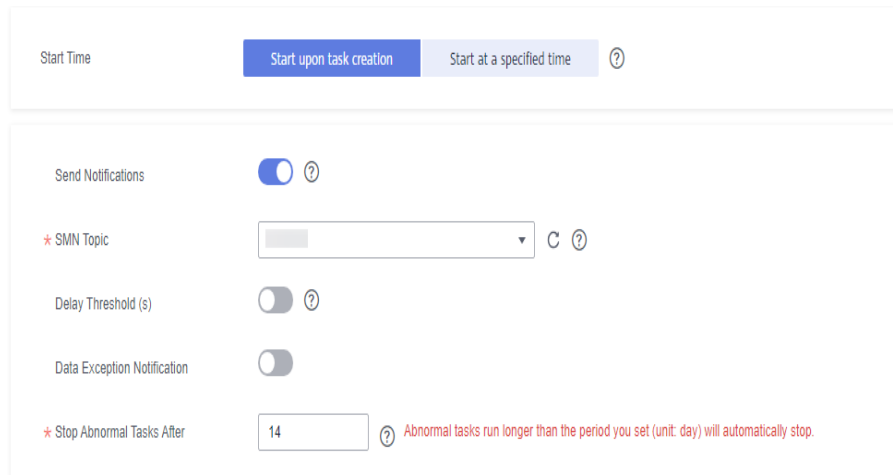



Table 5-211 Task startup settings

| Parameter | Description |
|------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |

| Parameter | Description |
|-----------------------------|--|
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none">• If the delay threshold is set to 0, no notifications will be sent to the recipient.• In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent.• Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none">• You can set this parameter only for pay-per-use tasks.• Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies

for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.20 From DB2 for LUW to GaussDB Distributed

Supported Source and Destination Databases

Table 5-212 Supported databases

| Source DB | Destination DB |
|---|---------------------|
| DB2 for LUW 9.7, 10.1, 10.5, 11.1, and 11.5 | GaussDB distributed |

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

- For more information about the impact of DRS on databases, see [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 and destination databases, [modify the connection information](#) in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 5-213 Precautions

| Type | Restrictions |
|----------------------|--|
| Database permissions | <ul style="list-style-type: none"> ● Source database <ul style="list-style-type: none"> - To perform full synchronization, the CONNECT and DATAACCESS permissions are required. - To perform full+incremental synchronization, the DBADM permission is required. - If you select full+incremental synchronization, ensure that the archive log of the source database is enabled and the Datacapture attribute of the table to be synchronized is Y. ● Destination database <ul style="list-style-type: none"> - Database-level permission: Log in to the postgres base database as user root or other DATABASE users with the sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE. - Schema-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user. - Table-level permission: Log in to the database as user root or user DATABASE with the sysadmin role, or the owner of the database, and grant the SELECT, UPDATE, INSERT, and DELETE permissions for all tables in the SCHEMA to the user. - If gs_loader is used to create system catalogs (such as public.pgxc_copy_error_log and public.gs_copy_summary) in the destination database, the all privilege permission is required for accessing the system catalogs. For details, see gs_loader. |

| Type | Restrictions |
|------------------------|---|
| Synchronization object | <ul style="list-style-type: none"> ● During full synchronization, tables, common indexes, primary 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 maximum precision supported by the timestamp type is 6. ● During full synchronization, the table structure does not support bitmap indexes, inverted indexes, function indexes, and XML indexes. ● The LOB type supported by incremental synchronization cannot exceed 10 MB. ● Tables whose default values contain expression functions cannot be synchronized. ● Temporary tables in the source database cannot be synchronized. ● DDL statements executed in the source database cannot be synchronized. ● For a full+incremental task, tables without primary keys that contain XML fields cannot be synchronized. ● If a floating point number is used as the primary key, incremental synchronization may cause data inconsistency. |
| Source database | <ul style="list-style-type: none"> ● Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _. DRS does not support non-ASCII characters or special characters .>`<`\, ?!". Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases. ● Only the GBK and UTF8 character sets can be synchronized. |

| Type | Restrictions |
|----------------------|--|
| Destination database | <ul style="list-style-type: none">• The destination database must be a GaussDB distributed instance.• Ensure that a database named in lowercase letters has been created in destination database.• Do not use foreign keys for tables during synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.• The destination table can contain more columns than the source table. However, the following failures must be avoided: Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. |

| Type | Restrictions |
|-------------|---|
| Precautions | <ul style="list-style-type: none"> ● Do not use an imprecise value type as the primary key in the database because the synchronization of UPDATE and DELETE statements in the DRS incremental scenario will be affected. ● ARM VMs are not supported. ● If a table that does not have a primary key contains LOB or LONG data types, data may be inconsistent during incremental synchronization. ● If the character sets of the source database are different from those of the destination database, adjust the field length of the destination database based on the site requirements. For example, the source character set is GBK and the destination character set is UTF8, a Chinese character of GBK occupies two bytes, and a Chinese character of UTF8 occupies three bytes. When CHAR or VARCHAR data type is synchronized to the destination database, the field length may exceed the defined length, so you need to increase the field length to 1.5 times of that in the source database. ● After a task is started, the detach operation on a partition table may cause data inconsistency. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During the synchronization, the source database cannot contain uncommitted DDL transactions. ● Full synchronization consists of two phases: table structure synchronization (including indexes) and real-time synchronization. If the structure of a table is created in the destination database, real-time synchronization starts. If a table fails to be synchronized, you can restart the task to synchronize the table data. However, the table structure is not synchronized, so you must manually create the table in the destination database. ● When the structure of a partitioned table is synchronized during a full synchronization, the partitioned table is converted to a non-partitioned table. ● During the full synchronization of table structures, only default value constraints of the character string or number type are supported. Default value constraints of the function or sequence type are not supported. If necessary, create default value constraints for the corresponding table in the destination database. ● Do not write data to the destination database table during full synchronization. Otherwise, data will be inconsistent. ● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated. ● You can add additional objects in the incremental synchronization phase. |

| Type | Restrictions |
|------|---|
| | <ul style="list-style-type: none"> During an incremental synchronization, if you update the primary key column or the first column of a table that does not have a primary key in the source database DB2 for LUW, the distribution column will be updated in the GaussDB database, which may cause data inconsistency. Thus, do not update the preceding columns. |

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-204 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-8131 [Text input]

Description: [Text area]

0/256

Table 5-214 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-205 Synchronization instance details

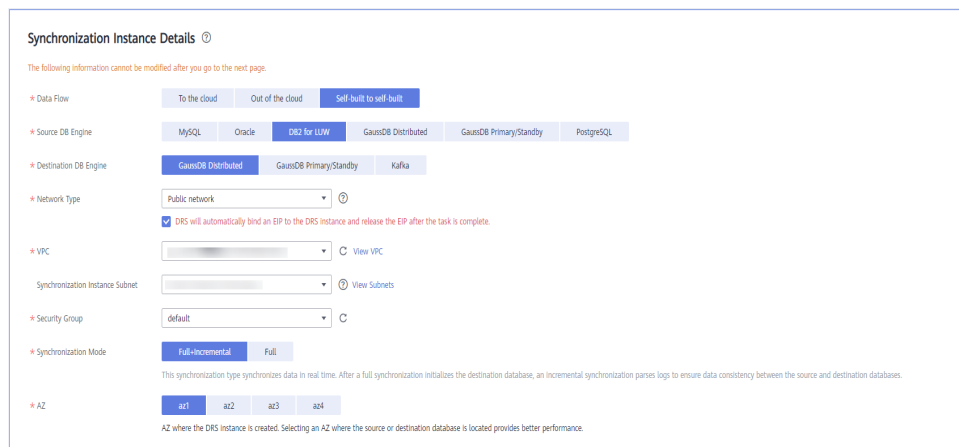


Table 5-215 Synchronization instance settings

| Parameter | Description |
|---------------------------------|---|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select DB2 for LUW . |
| Destination DB Engine | Select GaussDB Distributed . |
| Network Type | Public network is used as an example. Available options: Public network and VPN or Direct Connect |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |

| Parameter | Description |
|----------------------|--|
| Synchronization Mode | <ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. - Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. |
| 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 Project and Tags

Figure 5-206 Enterprise projects and tags

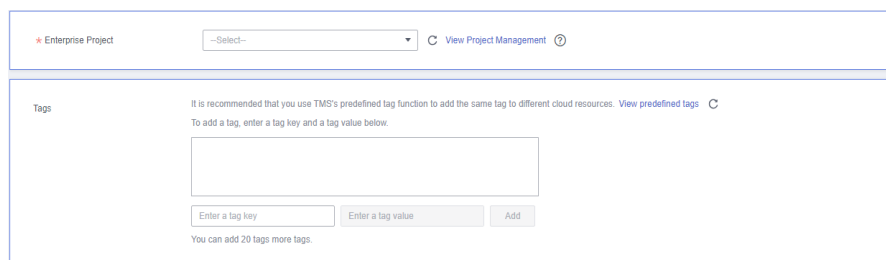


Table 5-216 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> – Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. – If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. – After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, click **Next**.

Figure 5-207 Source database information

Source Database

System databases, users, parameters, and jobs will not be migrated. You need to manually import users and jobs to the destination database and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

Database Name

SSL Connection

Synchronize Driver Synchronized.jcc-11.5.6.0(2).jar

Table 5-217 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. |


| Parameter | Description |
|--------------------|--|
| Database Password | The password for the source database username. |
| Database Name | The name of the database to which the synchronization objects belong in the source DB2 for LUW. |
| SSL Connection | <p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. Currently, only DB2 for LUW 10.5 and later versions support SSL. DB2 for LUW 9.7 and 10.1 do not support SSL.</p> <p>NOTE</p> <ul style="list-style-type: none"> The maximum size of a single certificate file that can be uploaded is 500 KB. If the SSL certificate is not used, your data may be at risk. Three certificate files need to be uploaded for the DB2 for LUW SSL connection: JKS, kdb, and sth DRS does not verify the kdb and sth certificates during the connection test. You need to ensure that the files are correct to prevent subsequent task failures. |
| Synchronize Driver | You need to manually add and upload the JDBC driver package corresponding to the source DB2 for LUW, and then test the connection. |

 **NOTE**

The IP address, username, and password of the source database are encrypted and stored in DRS and will be cleared after the task is deleted.


Figure 5-208 Destination database information

Destination Database

IP Address or Domain Name 

Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password 

This button is available only after the replication instance is created successfully.

Table 5-218 Destination database settings

| Parameter | Description |
|---------------------------|---|
| IP Address or Domain Name | IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535 You can enter up to 10 groups of IP addresses or domain names of the destination database. Separate multiple values with commas (.). Example: 192.168.0.1:8000,192.168.0.2:8000 |
| Database Username | The username for accessing the destination database. |
| Database Password | The password for the database username. |

 **NOTE**

The username and password of the destination database are encrypted and stored in DRS, and will be cleared after the task is deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization object, enter the name of the destination database, and click **Next**.

Figure 5-209 Synchronization mode

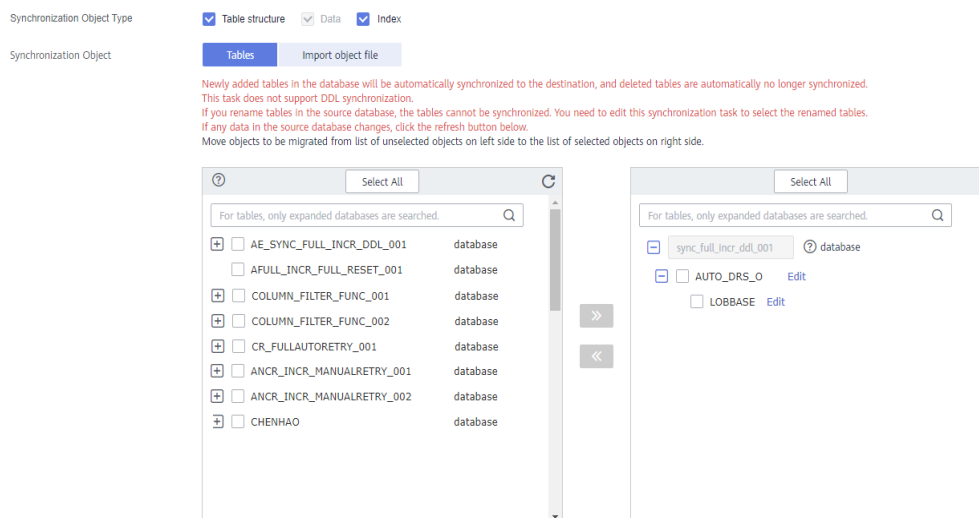



Table 5-219 Synchronization Object

| Parameter | Description |
|-----------------------------|---|
| Synchronization Object Type | <p>You can select Table structure, Data, or Index for Synchronization Object Type for full synchronization.</p> <ul style="list-style-type: none"> • Data is selected by default. • If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized. • If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures. |
| 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 an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If you need to filter out data, set the corresponding rules by referring to "Filtering Data" in [Processing Data](#).

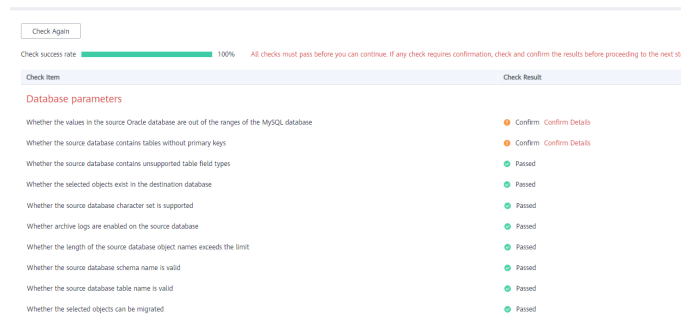
Step 6 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-210 Pre-check



NOTE

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 7 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 5-211 Task startup settings

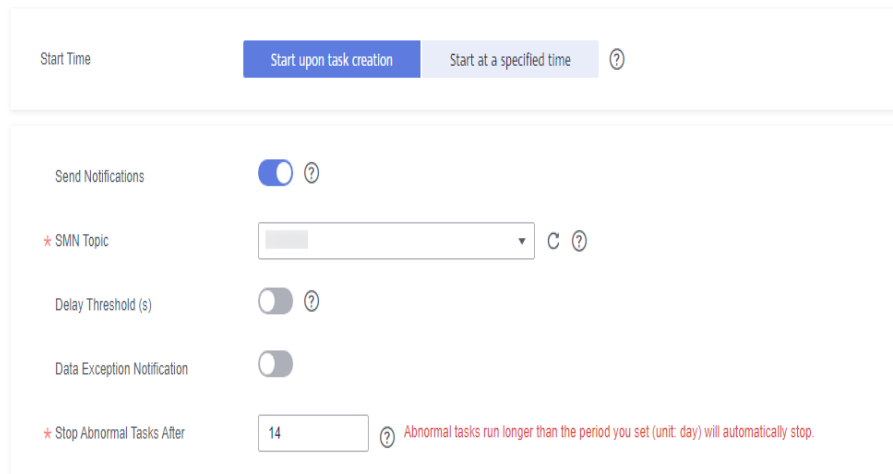



Table 5-220 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |

| Parameter | Description |
|-----------------------------|--|
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> If the delay threshold is set to 0, no notifications will be sent to the recipient. In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> You can set this parameter only for pay-per-use tasks. Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

5.21 From Microsoft SQL Server to Kafka

Supported Source and Destination Databases

Table 5-221 Supported databases

| Source DB | Destination DB |
|---|---------------------|
| <ul style="list-style-type: none"> On-premises Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) ECS-hosted Microsoft SQL Server (Enterprise Edition 2012, 2014, 2016, 2017 and 2019 and Standard Edition 2016 SP2 or later, 2017, and 2019) | Kafka 0.11 or later |

 **NOTE**

Only whitelisted users can use this function.

Supported Synchronization Objects

Table 5-222 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 5-222 Supported synchronization objects

| Type | Precautions |
|---------|---|
| Objects | <ul style="list-style-type: none"> Supported field types TINYINT, SMALLINT, INT, BIGINT, DECIMAL, NUMERIC, FLOAT, REAL, SMALLMONEY, MONEY, BIT, DATE, DATETIME, DATETIME2, DATETIMEOFFSET, TIME, TIMESTAMP, XML, CHAR, VARCHAR, NCHAR, NVARCHAR, BINARY, VARBINARY, IMAGE, HIERARCHYID, NTEXT, TEXT, and UNIQUEIDENTIFIER Unsupported field types: SQL_VARIANT, GEOMETRY, and GEOGRAPHY Scope of incremental synchronization <ul style="list-style-type: none"> DML statements, including INSERT, UPDATE, and DELETE, are supported. DDL statements are not supported. |

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 5-223](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 5-223 Database account permission

| Type | Incremental |
|----------------------|--|
| Source database user | Sysadmin or view server state and db_datareader or db_owner permissions for databases to be synchronized |

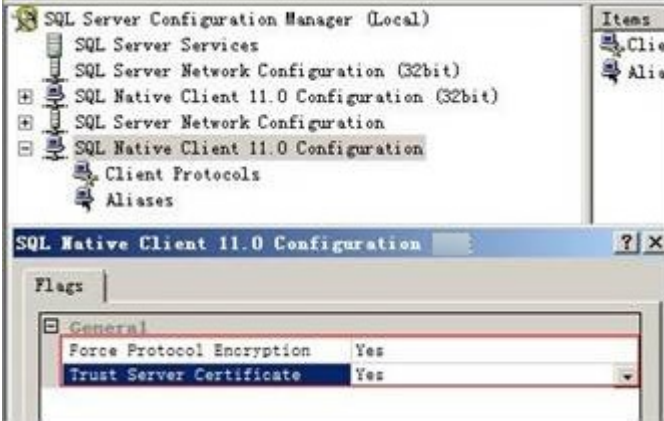
Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.
- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

DRS incremental synchronization consists of three phases: task start, incremental synchronization, and task completion. To ensure smooth synchronization, read the following notes before creating a synchronization task.

Table 5-224 Precautions

| Type | Restrictions |
|-----------------|---|
| Starting a task | <ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - The source database mode must be set to FULL. - The SQL Server Agent proxy service must be enabled for the source database. - If Force Protocol Encryption is set to Yes for the source database, Trust Server Certificate also must be set to Yes, as shown in Figure 5-212. <p style="text-align: center;">Figure 5-212 Client configuration</p>  <ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - If the source database contains disabled clustered indexes of tables, the synchronization fails. - The source database cannot contain the username cdc or schema. - Names of the databases, schemas, and tables to be synchronized in the source database can contain a maximum of 64 characters, including only letters, digits, underscores (_), and hyphens (-). - Names of the columns in the source table cannot contain the following special characters: []? ● Destination database parameter requirements: <ul style="list-style-type: none"> - The destination database is a Kafka database. - You are advised to set auto.create.topics.enable of Kafka to false. ● Other notes: <ul style="list-style-type: none"> - Do not perform primary/standby switchover on the source database. Otherwise, the synchronization task will fail. - The source Microsoft SQL Server database using TLS 1.0 or TLS 1.1 cannot be synchronized. If data synchronization is required, you are advised to upgrade TLS version of the source database to TLS 1.2 or later. |

| Type | Restrictions |
|-----------------------------|--|
| | <ul style="list-style-type: none"> – Incremental synchronization of the source SQL Server database is based on the CDC capability provided by the SQL Server database. If CDC of the source SQL Server database is abnormal (for example, CDC is disabled or the transaction log is full), incremental synchronization will be affected. |
| Incremental synchronization | <ul style="list-style-type: none"> • Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail. • Do not write data to the destination database. Otherwise, data may be inconsistent. • Do not delete the topic for receiving DRS data in Kafka. Otherwise, the task may fail. • DDL operations performed on the source database will not be synchronized to the destination database. • The IMAGE, TEXT, and NTEXT big data types cannot be deleted. • You can add additional synchronization objects. |

Procedure

This section uses self-managed Microsoft SQL Server to Kafka synchronization as an example to describe how to use DRS to configure a real-time synchronization task over a public network.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, select a region and project, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Figure 5-213 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 includes:

- 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 5-225 Task information

| Parameter | Description |
|-------------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance details

Figure 5-214 Synchronization instance details

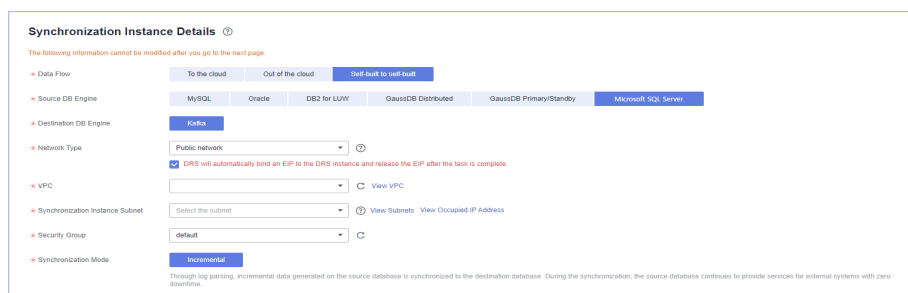


Table 5-226 Synchronization instance settings

| Parameter | Description |
|-----------------------|--|
| Data Flow | Choose Self-built to self-built . |
| Source DB Engine | Select Microsoft SQL Server . |
| Destination DB Engine | Select Kafka . |

| Parameter | Description |
|---------------------------------|---|
| Network Type | Available options: VPC , Public network and VPN or Direct Connect . Public network is used as an example. <ul style="list-style-type: none"> - VPC is suitable for data synchronization between cloud databases of the same account in the same region. - Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. - VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |
| VPC | Select an available VPC. |
| Synchronization Instance Subnet | Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed. |
| Security Group | Select a security group. You can use security group rules to allow or deny access to the instance. |
| Synchronization Mode | Incremental: Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- AZ

Figure 5-215 AZ




Table 5-227 Task AZ

| Parameter | Description |
|-----------|---|
| AZ | Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance. |

- Enterprise Project and Tags

Figure 5-216 Enterprise projects and tags




Table 5-228 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |
| Tags | <ul style="list-style-type: none"> Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 On the **Configure Source and Destination Databases** page, wait until the synchronization instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance. For details, see [Network Preparations](#).
- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

Figure 5-217 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 5-229 Source database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the source Microsoft SQL Server database. |
| Port | The port of the source database. Range: 1 – 65535 |
| Database Username | The username for logging in to the source Microsoft SQL Server database. |
| Database Password | The password for the database username. |

Figure 5-218 Destination database information

Destination Database

IP Address or Domain Name

Ensure that the entered addresses belong to the same DB instance.

Security Protocol

Table 5-230 Destination database settings

| Parameter | Description |
|---------------------------|--|
| IP Address or Domain Name | The IP address or domain name of the destination database. |
| Security Protocol | Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication . |

NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization policy and synchronization object, and click **Next**.

Figure 5-219 Synchronization Object

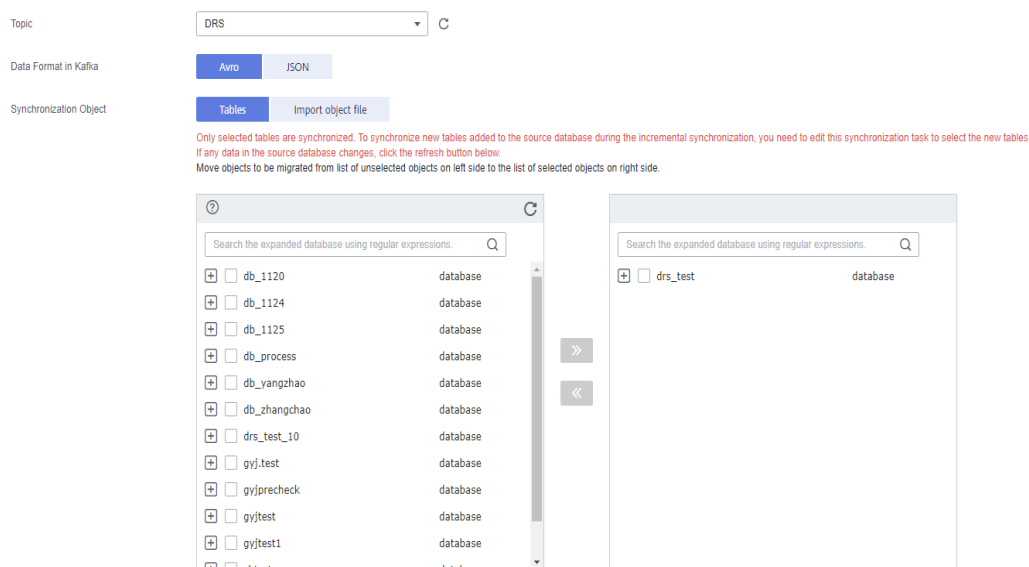



Table 5-231 Synchronization mode and object

| Parameter | Description |
|----------------------|--|
| Topic | Select the topic to be synchronized to the destination database. |
| Data Format in Kafka | Select the format of data delivered to Kafka. <ul style="list-style-type: none"> ● Avro refers to binary encoded format. ● Json refers to data interchange format. For details, see Kafka Message Format . |

| Parameter | Description |
|------------------------|--|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Import object file for Synchronization Object as required.</p> <ul style="list-style-type: none"> If you select Import object file for Synchronization Object, different tables can be synchronized to different topics at the destination end. For details about the import procedure and description, Importing Synchronization Objects. When you select Import object file, you can use the mapping function in Mapping Object Names only when the topic synchronization policy is set to A specific topic. Otherwise, topics are generated based on the name format. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

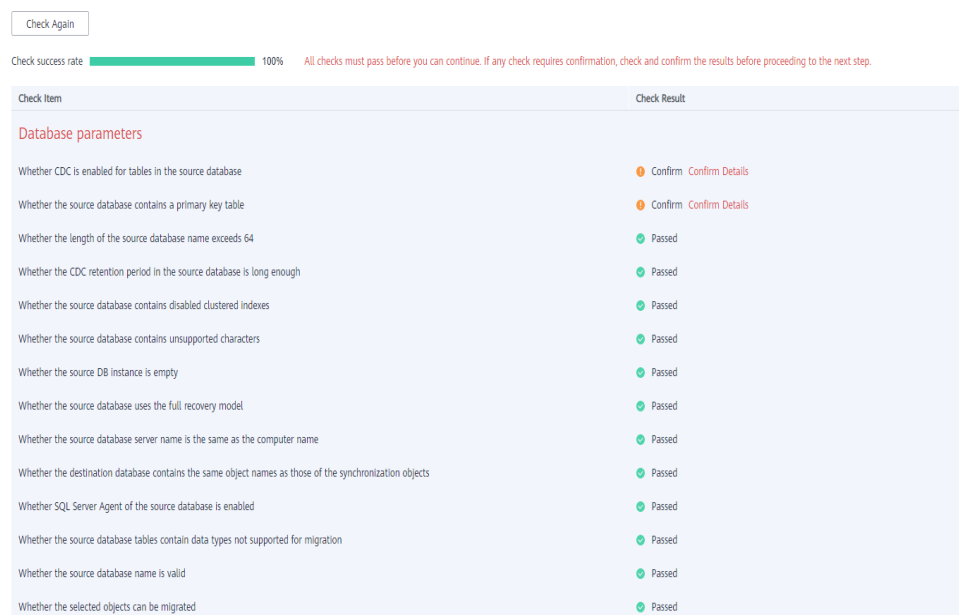
Step 5 On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 5-220 Pre-check



 **NOTE**

You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

Step 6 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.


Figure 5-221 Task startup settings

Table 5-232 Task startup settings

| Parameter | Description |
|--------------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |

| Parameter | Description |
|-----------------------------|---|
| Delay Threshold (s) | <p>During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database.</p> <p>If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | <p>This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal.</p> |
| Stop Abnormal Tasks After | <p>Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14.</p> <p>NOTE</p> <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 7 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies for resources for the task again. In this case, the IP address of the DRS instance changes.

----End

6 Task Management

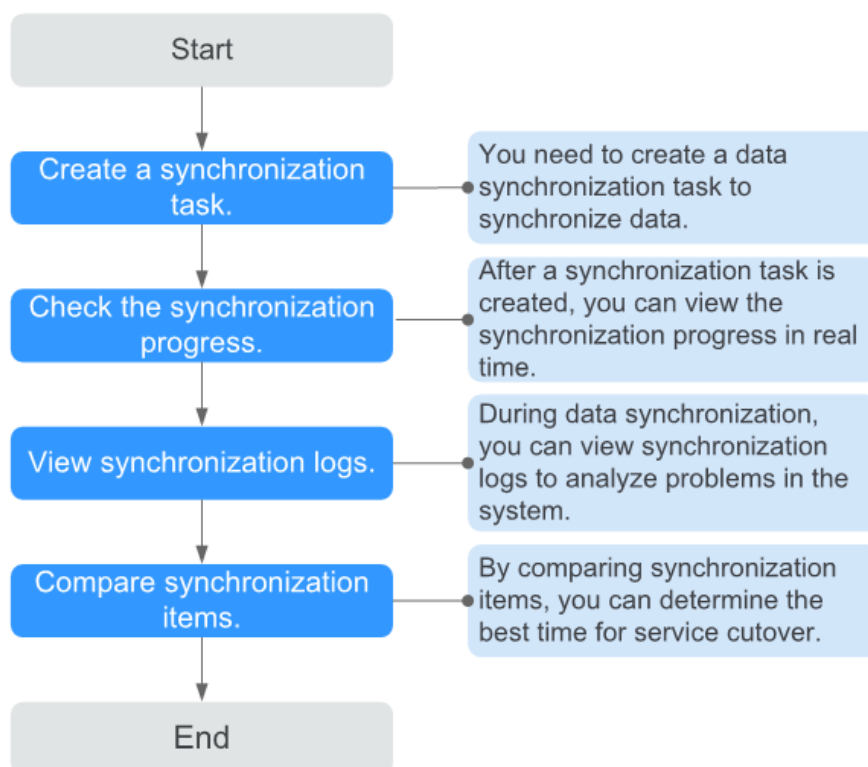
6.1 Creating a Synchronization Task

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.

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

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 text area with a height of 0/256 characters and a help icon.

Table 6-1 Task information

| Parameter | Description |
|-----------|--|
| Region | The region where the replication instance is deployed. You can change the region. |
| Project | The project corresponds to the current region and can be changed. |
| Task Name | The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_). |

| Parameter | Description |
|-------------|--|
| Description | The description consists of a maximum of 256 characters and cannot contain special characters !=<>'&"\ |

- Synchronization instance information

Figure 6-3 Synchronization instance information

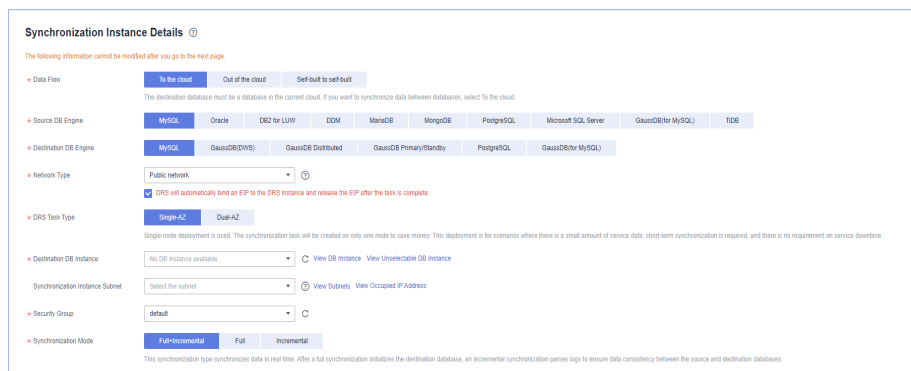


Table 6-2 Synchronization instance settings

| 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 | Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for data synchronization between cloud databases of the same account in the same region. – Public network is suitable for data synchronization from on-premises or external cloud databases to the destination databases bound with an EIP. – VPN or Direct Connect is suitable for synchronization of data between on-premises databases and cloud databases, between cloud databases of different accounts in the same region, or between cloud databases across regions. |

| Parameter | Description |
|---------------------------------|---|
| DRS Task Type | <p>Type of the DRS task. The value can be Single-AZ or Dual-AZ.</p> <ul style="list-style-type: none"> - Dual-AZ: This architecture provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization. This deployment is for scenarios where there is a lot of service data, long-term synchronization is required, and there are strict limits on how much service downtime can be tolerated. - Single-AZ: Single-node deployment is used. The synchronization task will be created on only one node to save money. This deployment is for scenarios where there is a small amount of service data, short-term synchronization is required, and there is no requirement on service downtime. <p>This option is available only in specific scenarios. For details, see Performing a Switchover for a Dual-AZ Task.</p> |
| Destination DB Instance | <p>The RDS DB instance you created.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The destination DB instance cannot be a read replica. - The source and destination DB instances can be the same DB instance. |
| Synchronization Instance Subnet | <p>Select the subnet where the synchronization instance is located. You can also click View Subnets to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p> |
| Security Group | <p>Select a security group. You can use security group rules to allow or deny access to the instance.</p> |

| Parameter | Description |
|----------------------|---|
| Synchronization Mode | <p>The synchronization mode supported by a DRS task. Full+Incremental is used as an example. For details about the underlying working principles for full or incremental synchronization, see Product Architecture and Function Principles.</p> <ul style="list-style-type: none"> – Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. – Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. – Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database. |

- Task Type

Figure 6-4 Task type





Table 6-3 Task type information

| Parameter | Description |
|----------------|---|
| Specifications | <p>DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization.</p> <p>NOTE DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be downgraded. For details, see Changing Specifications.</p> |

| Parameter | Description |
|-----------|---|
| AZ | <p>Select the AZ where you want to create the DRS task. Selecting the one housing the source or destination database can provide better performance.</p> <p>If DRS Task Type is set to Dual-AZ, you can specify Primary AZ and Standby AZ.</p> <p>Figure 6-5 AZ</p>  |

- Enterprise Project and Tags

Figure 6-6 Enterprise projects and tags




Table 6-4 Enterprise Project and Tags

| Parameter | Description |
|--------------------|--|
| Enterprise Project | <p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise project, see Enterprise Management User Guide.</p> <p>To customize an enterprise project, click Enterprise in the upper right corner of the console. The Enterprise Project Management Service page is displayed. For details, see Creating an Enterprise Project in <i>Enterprise Management User Guide</i>.</p> |

| Parameter | Description |
|-----------|---|
| Tags | <ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 20 tags. - If your organization has configured tag policies for DRS, add tags to tasks based on the policies. If a tag does not comply with the policies, task creation may fail. Contact your organization administrator to learn more about tag policies. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management. |

 **NOTE**

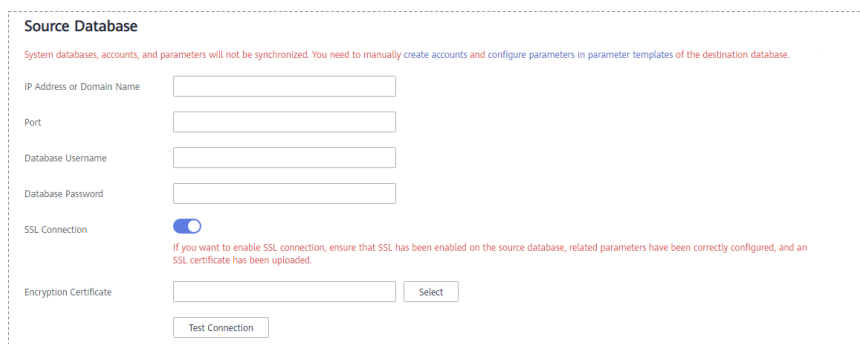
If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, click **Test Connection** for both the source and destination databases to check whether they have been connected to the synchronization instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

In different data flow scenarios, the source and destination database settings are different. Specify the required parameters based on the GUI.

- 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. |

| Parameter | Description |
|-------------------|---|
| 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 disabled, your data may be at risk. |

 **NOTE**

The IP address, port, username, and password of the source database are encrypted and stored in the database and the synchronization instance, and will be cleared after the task is deleted.


- Destination database information

Figure 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. |

| Parameter | Description |
|-------------------|---|
| 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 disabled, your data may be at risk. |

 **NOTE**

The username and password of the destination database are encrypted and stored in the database and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 6-9 Synchronization mode

Flow Control: Yes No ⓘ

Incremental Conflict Policy: Ignore Overwrite ⓘ
If synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database, the existing data will be overwritten.

Filter DROP DATABASE: Yes No

Synchronize: Normal index Incremental DDLs ⓘ

Data Synchronization Topology: One-to-one One-to-many Many-to-one ⓘ

Synchronize DDLs: Default Custom ⓘ
During database-level synchronization, all DDL operations in the binlog related to database objects, except DDL related those to permissions, are synchronized to the destination. Common DDL statements are CREATE_TABLE and RENAME_TABLE. During table-level synchronization, only DDL operations in the binlog related to the selected tables are synchronized. Common DDL statements are ADD_COLUMN, MODIFY_COLUMN, and ALTER_COLUMN.

| | | | | |
|--|---|--|---|---|
| <input checked="" type="checkbox"/> Add | <input checked="" type="checkbox"/> CREATE_TABLE | <input checked="" type="checkbox"/> ADD_COLUMN | <input checked="" type="checkbox"/> ADD_INDEX | <input checked="" type="checkbox"/> CREATE_INDEX |
| <input checked="" type="checkbox"/> Modify | <input checked="" type="checkbox"/> ADD_PRIMARY_KEY | <input checked="" type="checkbox"/> RENAME_TABLE | <input checked="" type="checkbox"/> MODIFY_COLUMN | <input checked="" type="checkbox"/> CHANGE_COLUMN |
| <input checked="" type="checkbox"/> Delete (High-risk) | <input checked="" type="checkbox"/> RENAME_INDEX | <input checked="" type="checkbox"/> DROP_COLUMN | <input checked="" type="checkbox"/> DROP_INDEX | <input checked="" type="checkbox"/> DROP_TABLE |
| | <input checked="" type="checkbox"/> TRUNCATE_TABLE | <input checked="" type="checkbox"/> DROP_PRIMARY_KEY | | |

Synchronization Object: Tables Databases Import object file

If any data in the source database changes, click the refresh button below.
Move objects to be migrated from list of unselected objects on left side to the list of selected objects on right side.

Select All

Search the expanded database using regular expressions:

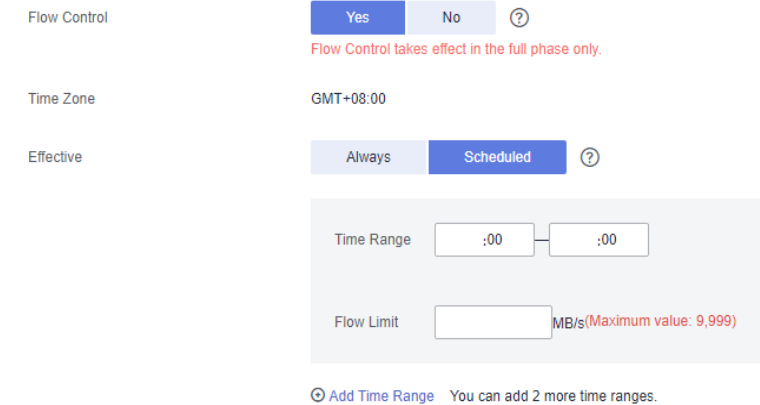
- alpha_60 database
- alpha_61 database
- create_database_escape_name_4 database
- db002 database
- db1 database
- db_specials2rds_vpc_null database
- mxz database
- parser_test_db database
- rows database
- test_db database
- test_db1 database
- test_db2 database

Select All

Search the expanded database using regular expressions:


- db database

Table 6-7 Synchronization mode and object

| Parameter | Description |
|--------------|---|
| Flow Control | <p>You can choose whether to control the flow. Flow Control takes effect in the full phase only.</p> <ul style="list-style-type: none"> Yes You can customize the maximum synchronization speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 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. If you select Ignore, data in the source database may be inconsistent with that in the destination database. ● Overwrite Conflicting data will be overwritten. |
| Filter DROP DATABASE | <p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none"> ● If you select Yes, the database deletion operation performed on the source database is not synchronized during data synchronization. ● If you select No, related operations are synchronized to the destination database during data synchronization. |
| Synchronize | <p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize 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 position after the start point during an incremental synchronization.</p> <p>Run show master status to obtain the start point of the source database and set File, Position, and Executed_Gtid_Set as prompted.</p> |
| 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> |

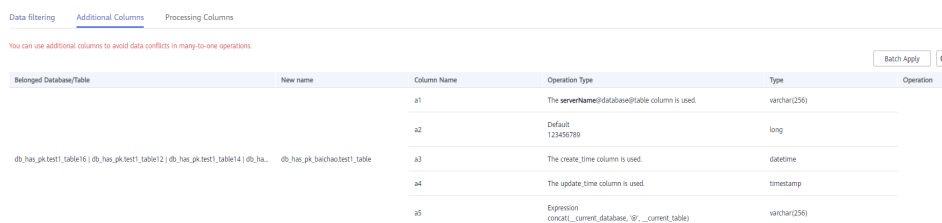
| Parameter | Description |
|------------------|---|
| Synchronize DDLs | <p>This parameter is available when Incremental DDLs is selected for Synchronize. Select DDL type for incremental synchronization. You can select Default or Custom based on your service requirements.</p> <ul style="list-style-type: 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. If Incremental DDLs is selected for Synchronize, but no DDL type is selected for Custom, DDLs will not be synchronized by default. <p>NOTE</p> <ul style="list-style-type: none"> ● Only whitelisted users can synchronize incremental DDL operations. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket. ● One-to-one and one-to-many scenarios: If the DDL usage of the source and destination databases must be consistent, high-risk DDLs must be synchronized. If you do not want a high-risk DDL to be performed in the destination, deselect the high-risk DDL to protect destination data. However, this may cause the synchronization to fail. However, filtering DDL may cause synchronization to fail, for example, column deletion. ● Many-to-one scenarios: Synchronize only the Add Column operation, or tasks may fail or data may be inconsistent due to changes in destination tables. |

| Parameter | Description |
|------------------------|---|
| Synchronization Object | <p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables, Import object file, or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one in the right pane. For details, see Mapping Object Names. <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. In the many-to-one scenario, if you want to change the saved database mapping name during table-level synchronization, you need to expand the database. For details about how to import an object file, see Importing Synchronization Objects. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces. |

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Columns**, or **Processing Columns**. For details about how to configure related rules, see [Processing Data](#).

Figure 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) | |

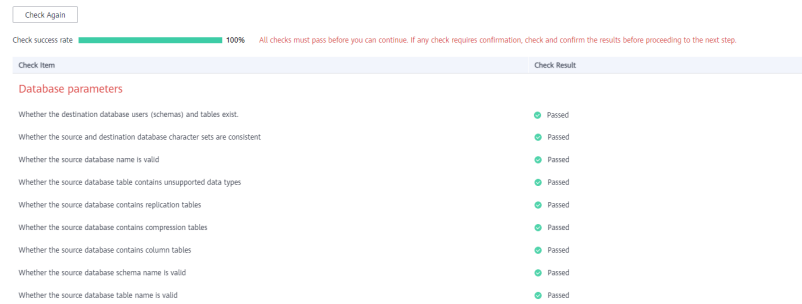
Step 6 On the **Check Task** page, check the synchronization task.

- 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 [Solutions to Failed Check Items](#) 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 Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After**, confirm that the configured information is correct, select the check box before the agreement, and click **Submit** to submit the task.

Figure 6-13 Task startup settings

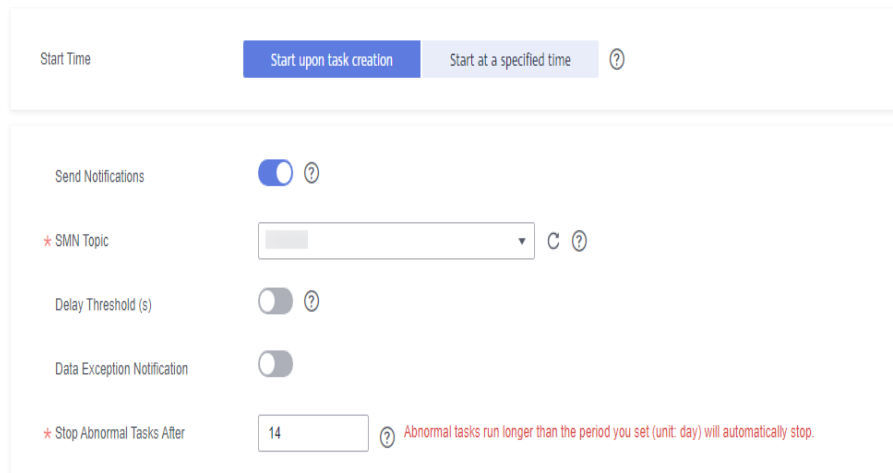



Table 6-8 Task startup settings

| Parameter | Description |
|------------|---|
| Start Time | Set Start Time to Start upon task creation or Start at a specified time based on site requirements. NOTE After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours. |

| Parameter | Description |
|-----------------------------|--|
| Send Notifications | This parameter is optional. After enabled, select a SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send a notification. |
| SMN Topic | This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see Simple Message Notification User Guide . |
| Delay Threshold (s) | During an incremental synchronization, a synchronization delay indicates a time difference (in seconds) of synchronization between the source and destination database. If the synchronization delay exceeds the threshold you specify, DRS will send alarms to the specified recipients. The value ranges from 0 to 3,600. To avoid repeated alarms caused by the fluctuation of delay, an alarm is sent only after the delay has exceeded the threshold for six minutes. NOTE <ul style="list-style-type: none"> • If the delay threshold is set to 0, no notifications will be sent to the recipient. • In the early stages of an incremental synchronization, the synchronization delay is long because a large quantity of data is awaiting synchronization. In this case, no notifications will be sent. • Before setting the delay threshold, enable Send Notifications. |
| Data Exception Notification | This parameter is optional. After enabled, DRS will send a notification if the task data is abnormal. |
| Stop Abnormal Tasks After | Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is 14 . NOTE <ul style="list-style-type: none"> • You can set this parameter only for pay-per-use tasks. • Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees. |

Step 8 After the task is submitted, you can view and [manage it](#) on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you configure the task again, DRS applies

for resources for the task again. In this case, the IP address of the DRS instance changes.

----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 the full synchronization, DRS displays the progress overview. You can view the synchronization progress of structures, data, and indexes. The synchronization progress is displayed based on the number of synchronized objects. When the progress reaches 100%, the synchronization is complete. The synchronization of data and indexes is relatively slow.
- During the full synchronization, you can view the synchronization progress of a specific object in the synchronization details. If the number of objects is the same as that of synchronized objects, the synchronization is complete. You can view the synchronization progress of each object in detail. During incremental synchronization, the progress details are not displayed. You can view the consistency status on the **Synchronization Comparison** tab.
- During an incremental synchronization, DRS displays the incremental synchronization delay. You can determine the synchronization status between the source and destination databases based on the delay. If the delay is 0, the source and destination databases are instantaneously consistent, and no new transaction needs to be synchronized.

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click **Synchronization Progress** to view table synchronization progress.
 - Check the synchronization progress in percentage.
 - When a full synchronization is complete, the progress reaches 100%.
 - After the full synchronization is complete, the incremental synchronization starts. You can view the incremental synchronization delay on the **Synchronization Progress** tab.
 - You can also view the incremental synchronization delay on the **Data Synchronization Management** page. When the incremental

synchronization delay exceeds the preset or default threshold, the value of the incremental synchronization delay is displayed in red in the task list.

- When the delay is 0s, data in the source database is synchronized to the destination database in real time.

 **NOTE**

"Delay" refers to the delay from when the transaction was submitted to the source database to when it is synchronized to the destination database and executed.

Transactions are synchronized as follows:

1. Data is extracted from the source database.
2. The data is transmitted over the network.
3. DRS parses the source logs.
4. The transaction is executed on the destination database.

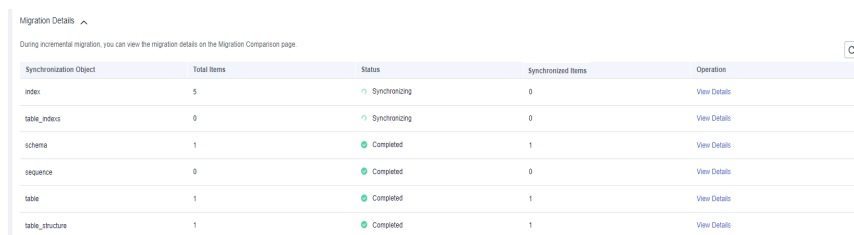
If the delay is 0, the source database is consistent with the destination database, and no new transactions need to be synchronized.




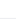
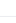

 **CAUTION**

Transactions that have been there for a long time without being submitted and frequent DDL operations may result in excessive synchronization delay.

- View the synchronization progress details. In the **Migration Details** area, locate the target synchronization object and click **View Details** in the **Operation** column to view the synchronization progress. After the incremental synchronization starts, the progress is not displayed. You can click the **Synchronization Comparison** tab to compare the data consistency.

Figure 6-14 Synchronization progress details




| Synchronization Object | Total Items | Status | Synchronized Items | Operation |
|------------------------|-------------|---|--------------------|------------------------------|
| index | 5 |  Synchronizing | 0 | View Details |
| table_index | 0 |  Synchronizing | 0 | View Details |
| schema | 1 |  Completed | 1 | View Details |
| sequence | 0 |  Completed | 0 | View Details |
| table | 1 |  Completed | 1 | View Details |
| table_structure | 1 |  Completed | 1 | View Details |

- 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**

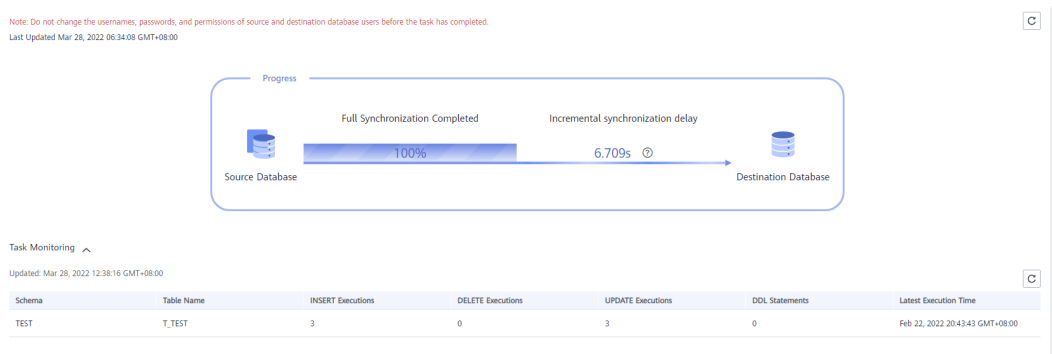
The displayed position information is updated every 10 seconds.

Figure 6-15 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 6-16 Task monitoring

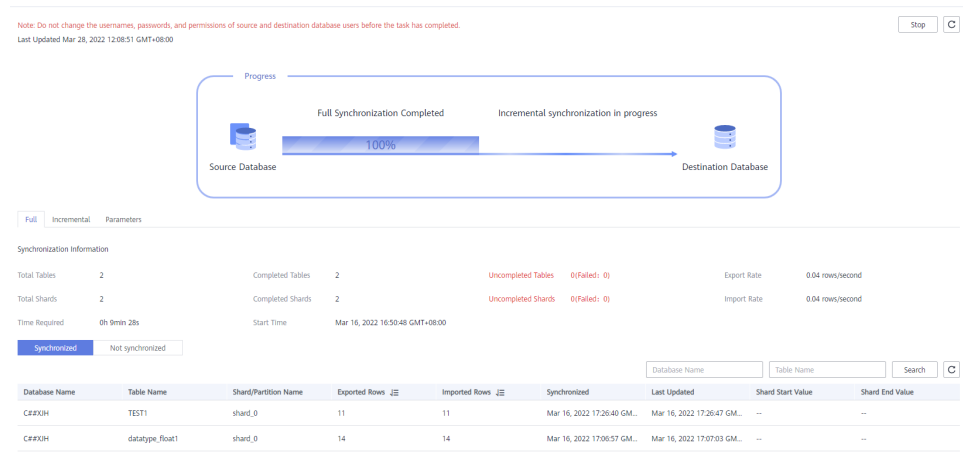


NOTE

1. After you perform the DML operation on the source database and run the commit command to make the operation take effect, the operation data can be displayed in the **Task Monitoring** list.
2. DRS collects statistics on the number of operations based on redo logs.
3. For tables with large object (LOB) data type columns, Oracle records more information in redo logs than the information generated by the actual operations that is performed. The INSERT and UPDATE operations are first performed on non-LOB columns and then on LOB columns. DRS collects statistics only from the redo logs, so the number of recorded operations may be inconsistent with the actual number of changed rows according to operation audit.
4. The Oracle MERGE statement can be converted into INSERT, UPDATE, and DELETE operations in the redo log. Thus, the number of rows changed by the MERGE statement are increased because the INSERT, UPDATE, and DELETE operations are separately recorded in the redo log.

Step 5 In the Oracle to GaussDB primary/standby or distributed, GaussDB primary/standby to MySQL, GaussDB distributed to MySQL, Oracle to MySQL (incremental), and Oracle to GaussDB(DWS) (incremental) synchronization scenarios, search for the detailed synchronization object information on the **Synchronization Progress** tab. The keywords for searching synchronization objects are case sensitive.

Figure 6-17 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 6-18 Synchronization logs



In addition, DRS can interconnect with Log Tank Service (LTS). After you enable log reporting to LTS, all logs generated by DRS instances will be uploaded to LTS for management. For details, see [Log Reporting](#).

----End

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 hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|-------------------|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| To the clou d | MySQL->MySQL | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | MySQL->PostgreSQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| To the clou d | MySQL -> GaussDB distributed | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | MySQL -> GaussDB primary/standby | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | MySQL->GaussDB(for MySQL) | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | MySQL->GaussDB(DWS) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | MySQL->MariaDB | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | PostgreSQL->PostgreSQL | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Suppo rted |
| To the clou d | PostgreSQL- >GaussDB(DWS) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | PostgreSQL -> GaussDB primary/standby | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | PostgreSQL -> GaussDB distributed | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--------------------------------------|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| To the clou d | DDM->MySQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | DDM->GaussDB(DWS) | Not suppor ted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | DDM->DDM | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | Oracle->MySQL | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | Oracle->GaussDB(for MySQL) | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | Oracle -> GaussDB primary/standby | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | Oracle -> GaussDB distributed | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | Oracle->DDM | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | Oracle->GaussDB(DWS) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| To the clou d | Oracle->PostgreSQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | DB2 for LUW -> GaussDB primary/standby | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | DB2 for LUW -> GaussDB distributed | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | DB2 for LUW- >GaussDB(DWS) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | TiDB->GaussDB(for MySQL) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | Microsoft SQL Server- >GaussDB(DWS) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | Microsoft SQL Server -> GaussDB primary/standby | Suppo rted | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted |
| To the clou d | Microsoft SQL Server -> GaussDB distributed | Suppo rted | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted |
| To the clou d | Microsoft SQL Server- >Microsoft SQL Server | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| To the clou d | MongoDB->DDS | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | MariaDB->MariaDB | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| To the clou d | MariaDB->MySQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | MariaDB->GaussDB(for MySQL) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| To the clou d | GaussDB(for MySQL)- >GaussDB(for MySQL) | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| From the clou d | MySQL->MySQL | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| From the clou d | MySQL->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | MySQL->CSS/ES | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | MySQL->MariaDB | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|---------------------------------------|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| From the clou d | DDM->MySQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | DDM->Oracle | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | DDM->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | DDS->MongoDB | Suppo rted | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted |
| From the clou d | DDS->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | PostgreSQL->PostgreSQL | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| From the clou d | PostgreSQL->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB primary/standby - > MySQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB primary/standby - > Oracle | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hroni zati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| From the clou d | GaussDB primary/standby - > Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB primary/standby - > GaussDB(DWS) | Suppo rted | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB primary/standby - > GaussDB distributed | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB primary/standby - > GaussDB primary/ standby | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB distributed -> MySQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB distributed -> Oracle | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB distributed -> GaussDB(DWS) | Suppo rted | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB distributed -> Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB distributed -> GaussDB distributed | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| From the clou d | GaussDB distributed- >GaussDB primary/standby | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB(for MySQL)- >MySQL | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| From the clou d | GaussDB(for MySQL)- >GaussDB(DWS) | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB(for MySQL)- >Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB(for MySQL)- >CSS/ES | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | GaussDB(for MySQL)- >Oracle | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| From the clou d | MariaDB->MariaDB | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| From the clou d | Microsoft SQL Server- >Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| Self- built -> Self- built | Oracle->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--------------------------------------|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| Self- built -> Self- built | Oracle -> GaussDB primary/standby | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| Self- built -> Self- built | Oracle -> GaussDB distributed | Suppo rted | Suppo rted | Suppo rted | Suppo rted | Not suppor ted |
| Self- built -> Self- built | MySQL->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| Self- built -> Self- built | MySQL->CSS/ES | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| Self- built -> Self- built | MySQL -> GaussDB primary/standby | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| Self- built -> Self- built | MySQL -> GaussDB distributed | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |
| Self- built -> Self- built | PostgreSQL->Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |
| Self- built -> Self- built | GaussDB primary/standby - > MySQL | Suppo rted | Suppo rted | Not suppor ted | Not suppor ted | Not suppor ted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|--|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| Self- built -> Self- built | GaussDB primary/standby - > Oracle | Suppo rted | Suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | GaussDB primary/standby - > Kafka | Not suppo rted | Not suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | GaussDB primary/standby - > GaussDB primary/ standby | Suppo rted | Suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | GaussDB distributed -> Oracle | Suppo rted | Suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | GaussDB distributed -> Kafka | Not suppo rted | Not suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | GaussDB distributed -> GaussDB distributed | Suppo rted | Suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | DB2 for LUW -> GaussDB primary/standby | Suppo rted | Suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |
| Self- built -> Self- built | DB2 for LUW -> GaussDB distributed | Suppo rted | Suppo rted | Not suppo rted | Not suppo rted | Not suppo rted |

| Sync hron izati on Dire ctio n | Data Flow | Objec t-level Comp arison | Row Comp arison | Value Comp arison | Dyna mic Comp arison | Accou nt- level Comp arison |
|--|---------------------------------|------------------------------------|-----------------------|-------------------------|-------------------------------|---|
| Self- built -> Self- built | Microsoft SQL Server- >Kafka | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted | Not suppor ted |

Constraints

- You can manually create a comparison task only when the task is in the incremental phase.
- During a comparison, the comparison items are case sensitive. If one of the source or destination database is case insensitive and the other one is case sensitive, the comparison result may be inconsistent.
- When a full migration task is complete, DRS automatically creates object-level and row comparison tasks. If operations are performed on data in the source database during data comparison, the comparison results may be inconsistent.
- If DDL operations were performed on the source database, you need to compare the objects again to ensure the accuracy of the comparison results.
- If data in the destination database is modified separately, the comparison results may be inconsistent.
- If the encoding of the source database character type is abnormal, the database driver will convert the character type to an abnormal code point during DRS migration or comparison. As a result, the values may be consistent but the bytes may be inconsistent.
- Currently, only tables with primary keys support value comparison. For tables that do not support value comparison, you can compare rows. Therefore, you can compare data by row or value based on scenarios.
- The DRS task cannot be suspended during value comparison. Otherwise, the comparison task may fail.
- Some data types do not support value comparison. For details, see [Which of the Following Data Types Are Not Supported By Value Comparison?](#)
- To prevent resources from being occupied for a long time, DRS limits the comparison duration. If the comparison duration exceeds the threshold, the comparison task stops automatically.
 - When a full migration task is complete, DRS automatically creates object-level and row comparison tasks. The comparison duration limits to 30 minutes. After the threshold, the comparison tasks automatically stop and the full migration task stops.

- For a row comparison task manually created in the incremental phase, if the source database is a relational database, the row comparison duration limits to 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration limits to 30 minutes.
- To avoid occupying resources, the comparison results of DRS tasks can be retained for a maximum of 60 days. After 60 days, the comparison results are automatically cleared.
- If you want to compare values for a DRS task, select large or ultra-large specifications when creating the DRS task.
- In the many-to-one row comparison scenario, the number of rows in the table in the source database is compared with that in the aggregation table mapped to the destination database.
- In the many-to-one synchronization scenario, value comparison is not recommended because data consistency cannot be ensured.
- Value comparison is not supported for a task in which tables in one database are mapped to multiple databases.
- 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.
- The empty character inserted into an Oracle database is processed as NULL. For tasks whose destination is an Oracle database, an empty string is considered as NULL. If data in the source database is empty and that in the destination database is NULL, the comparison result is consistent.
- During value comparison for synchronization from Oracle to GaussDB Distributed, if the LOB comparison policy is set to **Compare length**, the BLOB comparison is ignored because BLOB data in the distributed GaussDB instance fails to be queried using **DBE_LOB.LOB_GET_LENGTH**.

Impact on Databases

- Object comparison: System tables of the source and destination databases are queried, occupying about 10 sessions. The database is not affected. However, if there are a large number of objects (for example, hundreds of thousands of tables), the database may be overloaded.
- Row comparison: The number of rows in the source and destination databases is queried, which occupies about 10 sessions. The SELECT COUNT statement does not affect the database. However, if a table contains a large amount of data (hundreds of millions of records), the database will be overloaded and the query results will be returned slowly.
- Value comparison: All data in the source and destination databases is queried, and each field is compared. The query pressure on the database leads to high I/O. The query speed is limited by the I/O and network bandwidth of the source and destination databases. Value comparison occupies one or two CPUs, and about 10 sessions.
- Account comparison: The accounts and permissions of the source and destination databases are queried, which does not affect the database.

Estimated Comparison Duration

- Object comparison: Generally, the comparison results are returned within several minutes based on the query performance of the source database. If the amount of data is large, the comparison may take dozens of minutes.
- Row comparison: The SELECT COUNT method is used. The query speed depends on the database performance.
- Value comparison: If the database workload is not heavy and the network is normal, the comparison speed is about 5 MB/s.
- Account comparison: The results are returned with the object-level comparison results. If the number of objects is small, the results are returned in several minutes.

Prerequisites

- You have logged in to the DRS console.
- A synchronization task has been started.

Creating a Comparison Task

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 Choose **Synchronization Comparison**.

Step 3 Compare synchronization items.


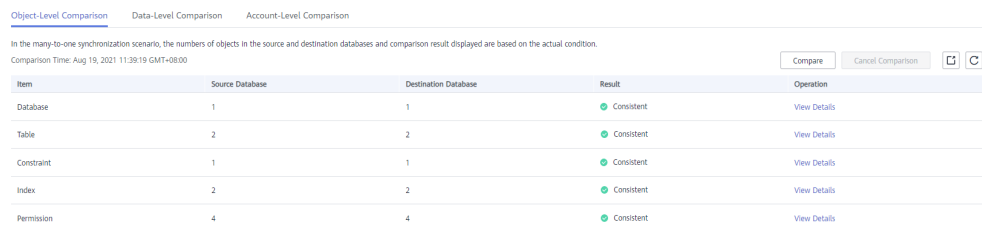
- On the **Object-Level Comparison** tab, click **Compare**. Wait for a while and click  to check whether the comparison results of the source and destination databases are consistent. Locate a comparison item you want to view and click **View Details** in the **Operation** column.

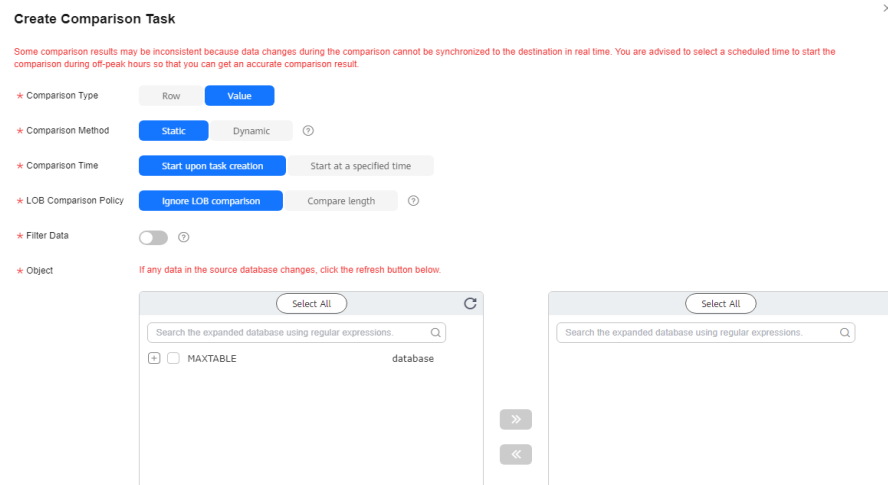
Figure 6-19 Object-level comparison



| 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 6-20 Creating a data-level comparison task



– **Comparison Type:** compares rows and values.

- Row comparison: checks whether the source table has the same number of rows as the destination table.

NOTE

- After a task enters the incremental comparison phase, you can create a row comparison task.
- For Oracle to GaussDB synchronization, the row comparison task is automatically triggered after the full synchronization is complete.

- Value comparison: checks whether the source table has the same data as the destination table.

NOTE

- After a task enters the incremental synchronization phase, you can create a value comparison task. After the full synchronization is complete, data in the source database cannot be changed. Otherwise, the comparison result will be inconsistent.
- After the synchronization from Oracle to GaussDB primary/standby and from Oracle to GaussDB distributed enters the incremental verification phase, the comparison service starts to extract data from the incremental logs of the tables involved in the task for continuous comparison. If you want to compare the status of other tables, you can cancel the ongoing comparison task and create a comparison task.

Value comparison only applies to tables with single-column primary key or unique index. You can use row comparison for tables that do not support value comparison. Therefore, you can compare data by row or value based on scenarios.

– **Comparison Policy:** DRS supports one-to-one and many-to-one comparison policies.

- **One-to-one:** compares the number of rows in a table in the source database with that in the table mapped to the destination database.

- **Many-to-one:** compares the number of rows in a table in the source database with that in the aggregate table mapped to the destination database.

 **NOTE**

If you select **Row Comparison** for **Comparison Type**, the **Comparison Policy** option becomes available.

- **Comparison Method:** DRS provides static and dynamic comparison methods.
 - **Static:** All data in the source and destination databases is compared. The comparison task ends as the comparison is completed. Static comparison can only be performed when there are no ongoing services.
 - **Dynamic:** All data in the source database is compared with that in the destination database. After the comparison task is complete, incremental data in the source and destination databases is compared in real time. A dynamic comparison can be performed when data is changing.

 **NOTE**

- If you select **Value** for **Comparison Type**, the **Comparison Method** option becomes available.
- During database-level synchronization, tables cannot be created in the source database during dynamic comparison. If you want to create a table in the source database, cancel the dynamic comparison first. After the new table is created and synchronized, restart the dynamic comparison.
- **Comparison Time:** You can select **Start upon task creation** or **Start at a specified time**. There is a slight difference in time between the source and destination databases during synchronization. Data inconsistency may occur. You are advised to compare migration items during off-peak hours for more accurate results.
- **LOB Comparison Policy:** You can select **Ignore LOB comparison** or **Compare length**.

 **NOTE**

LOB comparison policy can be set only for data synchronization from Oracle to GaussDB.

- **Ignore LOB comparison:** The system ignores LOB data during value comparison.
- **Compare length:** The system compares length of LOB data during value comparison.

You are advised to select **Ignore LOB comparison** because comparing length of LOB data increases the database load.
- **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 6-21 Account-level comparison

| Source Database Account Attribute | Source Database Account Name | Destination Database Account Attribute | Destination Database Account Name | Migration Comparison Time | Result |
|-----------------------------------|--|--|--|-------------------------------|------------|
| -- | ghuser1 | -- | ghuser1 | 2021-07-29 16:52:43 GMT+08:00 | Consistent |
| -- | ghlonguser567890qwertyuiopasdfghjklzxcvbnm | -- | ghlonguser567890qwertyuiopasdfghjklzxcvbnm | 2021-07-29 16:52:43 GMT+08:00 | Consistent |
| -- | ghuser2 | -- | ghuser2 | 2021-07-29 16:52:43 GMT+08:00 | Consistent |
| -- | ghlonguser567890qwertyuiopasdfghjklzxcvbnm | -- | ghlonguser567890qwertyuiopasdfghjklzxcvbnm | 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 6-22 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 6-23 Row comparison details

| Source Database | Destination Database | Result | Operation |
|-----------------|----------------------|------------|------------------------------|
| lf_test | lf_target2 | Consistent | View Details |

| Source Database Table Name | Destination Database Table Name | Source Database Table Rows | Destination Database Table Rows | Row Results | Row Differences |
|----------------------------|---------------------------------|----------------------------|---------------------------------|-------------|-----------------|
| lf_test | lf_target2 | 2 | 2 | Consistent | 0 |

Figure 6-24 Value comparison details

| Source Database | Destination Database | Result | Operation |
|-----------------|----------------------|------------|------------------------------|
| lf_test | lf_target2 | Consistent | View Details |

| Source Database Table Name | Destination Database Table Name | Value | Operation |
|----------------------------|---------------------------------|------------|-----------|
| lf_test | lf_target2 | Consistent | |

Tables Not Compared lf_test - lf_target2

If you need to compare tables that have been identified as incompatible, rectify the fault based on the comparison failure cause given in Value and then restart a new comparison task.

| Source Database Table Name | Destination Database Table Name | Value |
|----------------------------|---------------------------------|-------|
| No data available | | |

 **NOTE**

- You can cancel a running task at any time and view the comparison report of a canceled comparison task.
- You can sort the row comparison results displayed on the current page in ascending or descending order based on the number of rows in the source database table or the destination database table.
- If a negative number is displayed in the **differences** column, the number of rows in the destination database table is greater than that in the source database table. If a positive number is displayed in the **differences** column, the number of rows in the source database table is greater than that in the destination database table.

----End

Periodic Comparison

Periodic comparison indicates that DRS periodically compares the number of rows in the source database table with those in the destination database table and displays the comparison results.

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 Click the **Synchronization Comparison** tab.

Step 3 Click the **Periodic Comparison** tab and click **Modify Comparison Policy** to modify the comparison policy.

 **NOTE**

Only MySQL-to-MySQL, MySQL-to-GaussDB(for MySQL), MySQL-to-GaussDB(DWS), and GaussDB(for MySQL)-to-MySQL synchronization tasks support periodic comparison.

Step 4 In the **Modify Comparison Policy** dialog box, enable periodic comparison, configure the comparison frequency and time, and click **Yes**.

 **NOTE**

- After periodic comparison is enabled, DRS compares the number of rows at the scheduled time. You can view the comparison results on the **Data-Level Comparison** tab.
- After periodic comparison is disabled, only historical comparison results can be viewed.
- Modifications to the comparison policy settings take effect from the next comparison and do not affect the on-going periodic comparison tasks.
- During periodic comparison, the source and destination databases will be read. Perform the comparison during off-peak hours.
- During periodic comparison, ultra-large tables (those with more than 100 million rows) are automatically filtered out. You can use data-level comparison to spot check such large tables. It is not recommended that these large tables be compared periodically.

Figure 6-25 Modify Comparison Policy

Modify Comparison Policy ✕

Status
View comparison results in Data-Level Comparison.

Comparison Frequency
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 interrupted, and the results of comparisons that have been completed can still be viewed.

Comparison Type Row Comparison

Comparison Policy One-to-one Many-to-one ?

* 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 (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 checked periodically.

----End

6.5 Managing Objects

6.5.1 Editing Synchronization Objects

This section describes how to edit synchronization objects in an incremental synchronization task.

- For a normal incremental task, you can edit synchronization objects by adding or deleting databases and tables to be synchronized.
- For a failed incremental task, you can edit synchronization objects by changing the objects to be synchronized. If an incremental synchronization fails due to incorrect synchronization objects, you can use this function to remove the databases or tables and submit the task again to restore the task.
- During synchronization object editing of a failed incremental task, do not resume or reset the synchronization task before it is started. Otherwise, the synchronization object cache will be cleared, and the synchronization objects fail to be edited.

For details about the synchronization scenarios where synchronization objects can be edited in the incremental phase, see [Real-time Synchronization Scenarios Where Synchronization Objects Can Be Edited](#).


Prerequisites

You have logged in to the DRS console.

Method 1

- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task and click **Edit** in the **Operation** column.

Figure 6-26 Editing synchronization objects



- Step 2** On the **Set Synchronization Task** page, change the objects to be synchronized and click **Next**.

- You can search the expanded database using regular expressions.
- If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed.
- The name of the selected synchronization object cannot contain spaces.

- Step 3** On the **Process Data** page, set rules for a new table by referring to [Processing Data](#).

NOTE

The processing rules for a synchronized table cannot be modified.

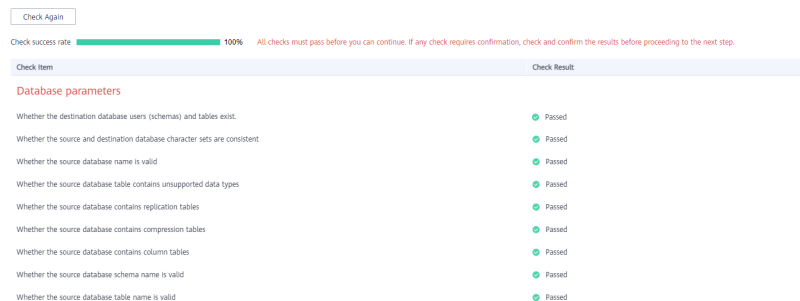
- Step 4** On the **Check Task** page, check the synchronization task.

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

For details about how to handle check failures, see [Solutions to Failed Check Items](#) in *Data Replication Service User Guide*.

- If all check items are successful, click **Next**.

Figure 6-27 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 5 On the **Confirm Task** page, specify **Start Time**, confirm that the configured information is correct, and click **Next**.

Step 6 Go back to the **Data Synchronization Management** page. In the synchronization task list, the current task status is **Incremental synchronization**, and a subtask in the **Modifying task** status is generated. 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 6](#) from method 1.

----End

Real-time Synchronization Scenarios Where Synchronization Objects Can Be Edited

To the cloud

- MySQL -> MySQL
- MySQL -> GaussDB(for MySQL)
- MySQL -> GaussDB Primary/Standby
- MySQL -> GaussDB Distributed
- MySQL->MariaDB

- DDM -> DDM
- DDM -> GaussDB(DWS)
- Oracle -> MySQL
- Oracle -> PostgreSQL
- Oracle -> GaussDB Primary/Standby
- Oracle -> GaussDB Distributed
- Oracle -> GaussDB(DWS)
- Oracle -> GaussDB(for MySQL)
- PostgreSQL -> PostgreSQL
- PostgreSQL -> GaussDB Distributed
- PostgreSQL -> GaussDB Primary/Standby
- MariaDB -> MariaDB
- MariaDB->MySQL
- MariaDB->GaussDB(for MySQL)
- DB2 for LUW -> GaussDB Distributed
- DB2 for LUW -> GaussDB Primary/Standby
- DB2 for LUW -> GaussDB(DWS)
- Microsoft SQL Server -> GaussDB(DWS)
- Microsoft SQL Server -> GaussDB Distributed
- Microsoft SQL Server -> GaussDB Primary/Standby
- GaussDB(for MySQL) -> GaussDB(for MySQL)

From the cloud

- MySQL -> MySQL
- MySQL -> CSS/ES
- MySQL -> Kafka
- MySQL -> Oracle
- MySQL->MariaDB
- DDM -> Kafka
- DDS -> Kafka
- PostgreSQL -> PostgreSQL
- PostgreSQL -> Kafka
- MariaDB -> MariaDB
- GaussDB(for MySQL) -> MySQL
- GaussDB(for MySQL) -> CSS/ES
- GaussDB(for MySQL) -> GaussDB(DWS)
- GaussDB(for MySQL) -> Oracle
- GaussDB(for MySQL) -> Kafka
- GaussDB Primary/Standby -> MySQL
- GaussDB Primary/Standby -> Oracle
- GaussDB Primary/Standby -> Kafka

- GaussDB Primary/Standby -> GaussDB(DWS)
- GaussDB Primary/Standby -> GaussDB Distributed
- GaussDB Primary/Standby -> GaussDB Primary/Standby
- GaussDB Distributed -> MySQL
- GaussDB Distributed -> Oracle
- GaussDB Distributed -> Kafka
- GaussDB Distributed -> GaussDB(DWS)
- GaussDB Distributed -> GaussDB Distributed
- GaussDB Distributed -> GaussDB Primary/Standby
- Microsoft SQL Server->Kafka

Self-built -> Self-built

- MySQL -> CSS/ES
- MySQL -> Kafka
- MySQL -> GaussDB Primary/Standby
- Oracle -> Kafka
- Oracle -> GaussDB Primary/Standby
- Oracle -> GaussDB Distributed
- PostgreSQL -> Kafka
- DB2 for LUW -> GaussDB Distributed
- DB2 for LUW -> GaussDB Primary/Standby
- GaussDB Primary/Standby -> Oracle
- GaussDB Primary/Standby -> Kafka
- GaussDB Primary/Standby -> GaussDB Primary/Standby
- GaussDB Distributed -> Oracle
- GaussDB Distributed -> Kafka
- GaussDB Distributed -> GaussDB Distributed
- Microsoft SQL Server->Kafka

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 6-28 Importing an object file



- Step 2** Click **Download Template**.
- Step 3** Download the template and enter information about the objects to be imported.
- Step 4** Click **Select File**. In the displayed dialog box, select the edited template.
- Step 5** Click **Upload**.

----End

6.5.3 Mapping Object Names

Data synchronization allows you to synchronize objects (including databases, schemas and tables) in a source 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.

This section describes how to map objects when configuring a data synchronization task. For details about the mapping relationship, see [Viewing Synchronization Mapping Information](#).

Precautions

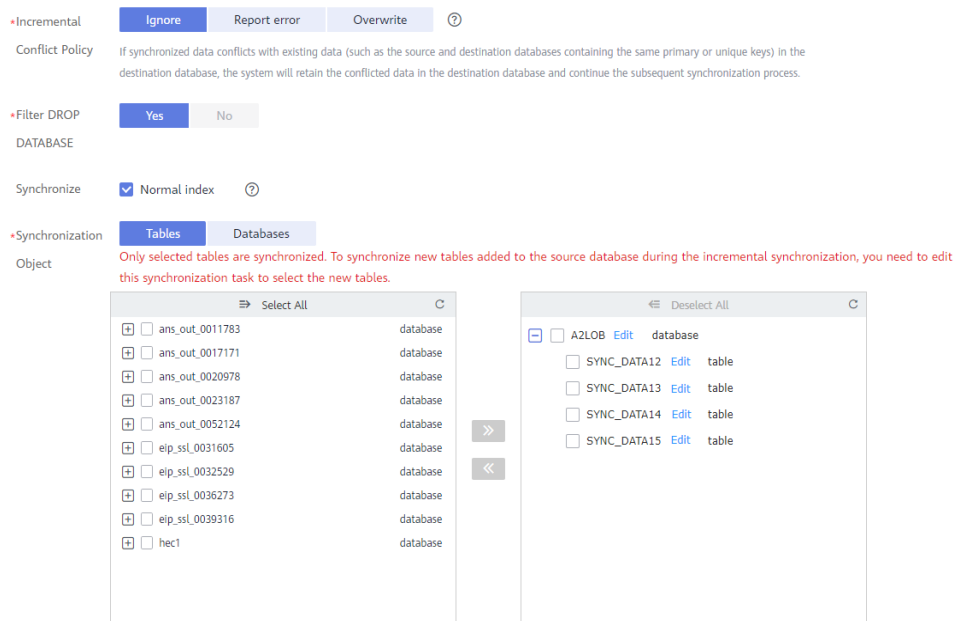
- Objects whose database names or table names contain newline characters cannot be mapped.
- If the destination DB is a type of PostgreSQL (for example, GaussDB(DWS), GaussDB, or PostgreSQL), the destination schema name cannot start with **pg_**. Otherwise, the migration fails.
- The following data flow types do not support schema mapping for tables of the smallserial, serial, and bigserial types:
 - GaussDB primary/standby -> GaussDB distributed
 - GaussDB primary/standby -> GaussDB primary/standby
 - GaussDB distributed -> GaussDB distributed
 - GaussDB distributed -> GaussDB primary/standby
- The following data flow types support many-to-one object mapping:
 - MySQL -> MySQL
 - MySQL -> GaussDB(for MySQL)
 - MySQL -> GaussDB(DWS)
 - MySQL -> Kafka
 - GaussDB(for MySQL) -> MySQL
 - GaussDB(for MySQL) -> GaussDB(for MySQL)
 - GaussDB(for MySQL) -> GaussDB(DWS)
 - Oracle -> GaussDB
 - Oracle -> GaussDB(DWS)

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. For example, when synchronizing database A in the source database to database B in the destination database, you need to map database name first.

- Step 1** On the **Set Synchronization Task** page, select the database that needs to be mapped from the synchronization objects on the right area and click **Edit**.

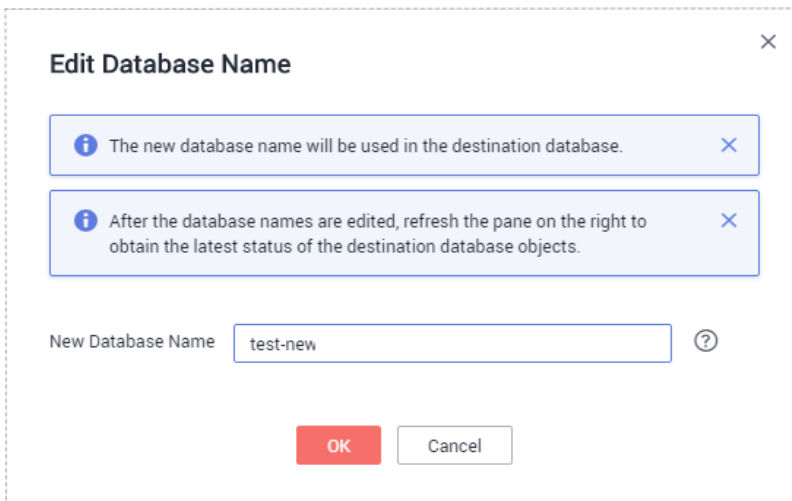
Figure 6-29 Mapping databases



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.

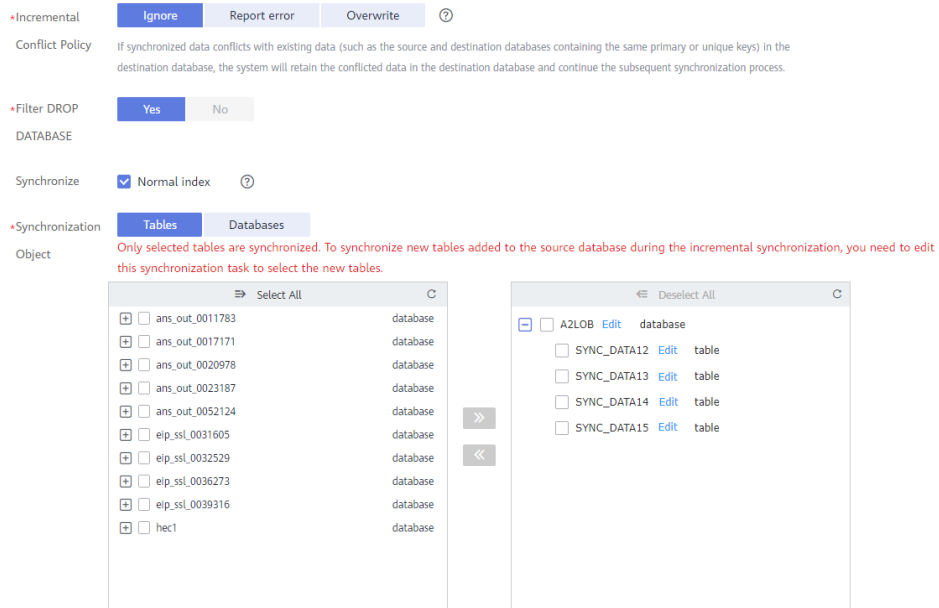
Figure 6-30 Changing a database name



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 6-31 Checking the result



----End

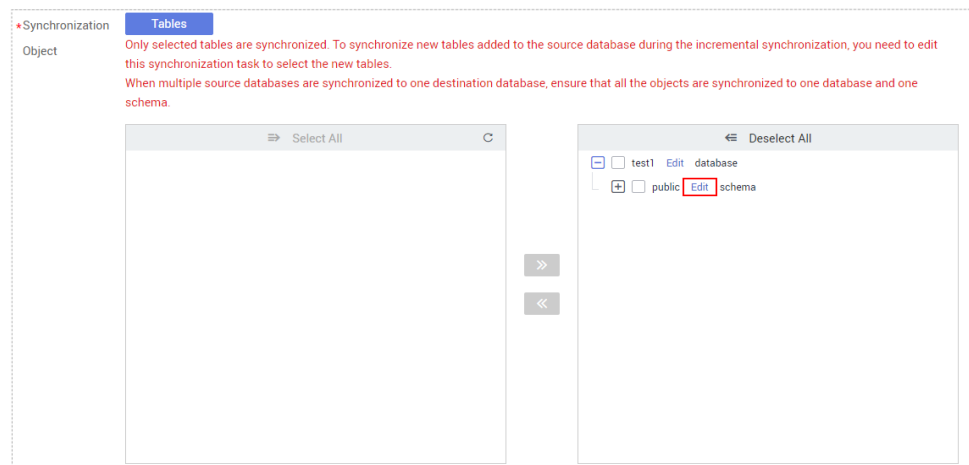
Mapping Schemas

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. For example, you need to synchronize schema A in the source database to schema B in the destination database.

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**.

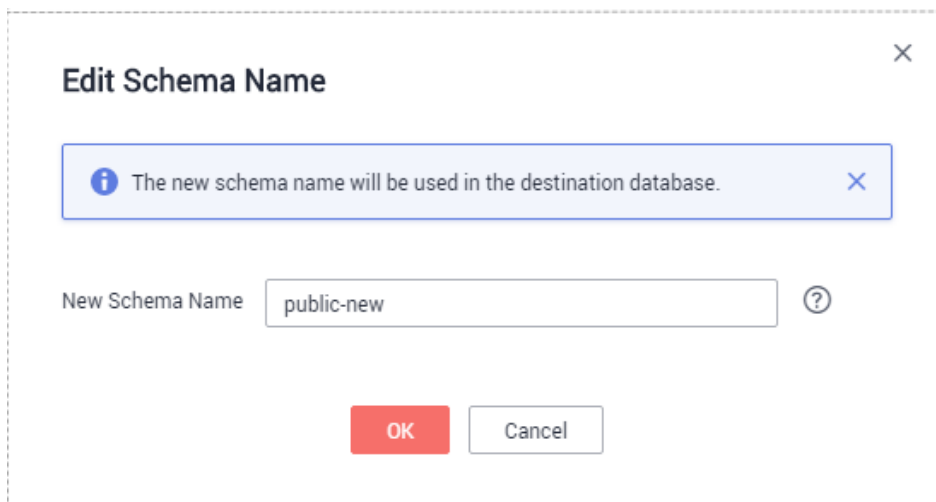
Figure 6-32 Mapping schemas



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.

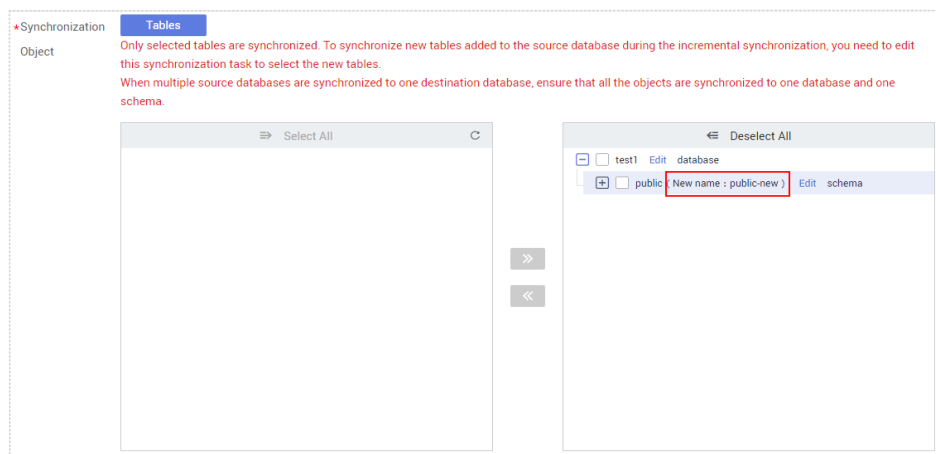
Figure 6-33 Edit Schema Name



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 6-34 Checking the result



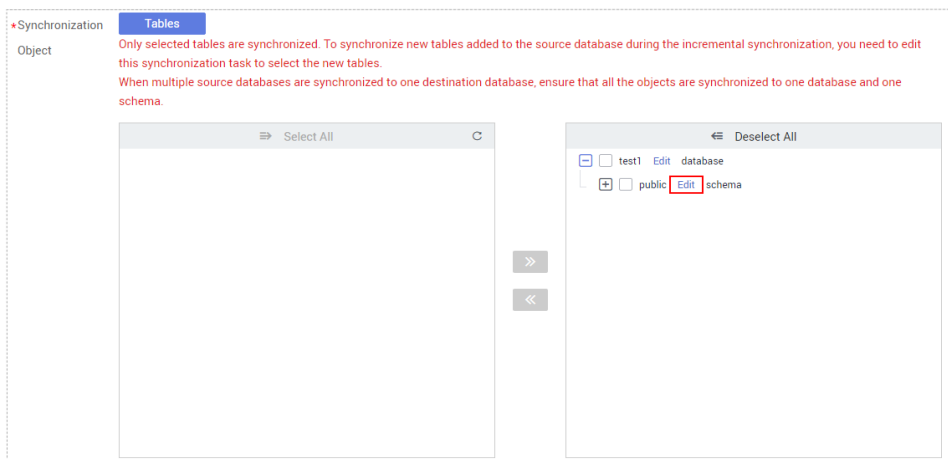
----End

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. For example, when synchronizing table A in the source database to table B in the destination database you need to map table names first.

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**.

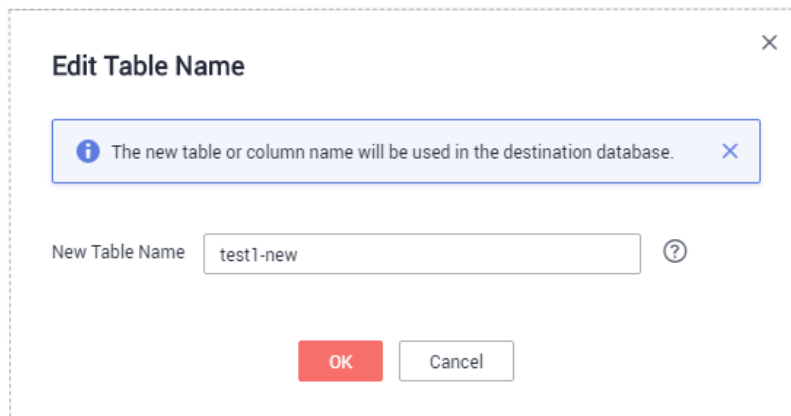
Figure 6-35 Table name mapping



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.

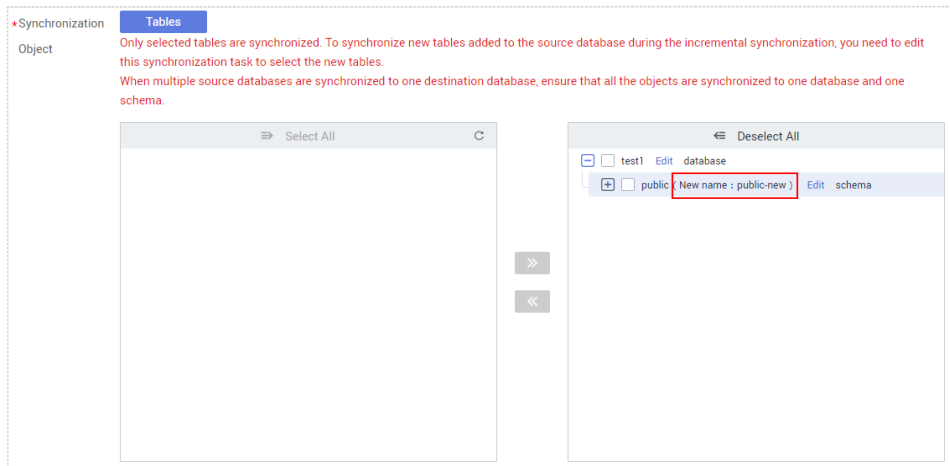
Figure 6-36 Changing a table name



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 6-37 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 6-38 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 hroni zatio n Direc tion | Data Flow | Data Filtering | Adition al Column | Column Processin g |
|--|---------------------------------------|-------------------|-------------------------|--------------------------|
| To the cloud | MySQL->MySQL | Supporte d | Supporte d | Supporte d |
| To the cloud | MySQL->GaussDB(DWS) | Supporte d | Supporte d | Not supported |
| To the cloud | MySQL->GaussDB(for MySQL) | Supporte d | Supporte d | Not supported |
| To the cloud | MySQL->MariaDB | Supporte d | Supporte d | Supporte d |
| To the cloud | DDM->GaussDB(DWS) | Not supported | Supporte d | Not supported |
| To the cloud | Oracle->GaussDB(DWS) | Supporte d | Supporte d | Not supported |
| To the cloud | Oracle->MySQL | Supporte d | Not supported | Not supported |
| To the cloud | Oracle->GaussDB(for MySQL) | Supporte d | Not supported | Not supported |
| To the cloud | Oracle -> GaussDB primary/ standby | Supporte d | Not supported | Supporte d |
| To the cloud | Oracle -> GaussDB distributed | Supporte d | Not supported | Supporte d |

| Sync hroni zatio n Direc tion | Data Flow | Data Filtering | Additional Column | Column Processin g |
|-------------------------------|--|----------------|-------------------|--------------------|
| To the cloud | DB2 for LUW -> GaussDB primary/standby | Supporte d | Not supported | Not supported |
| To the cloud | DB2 for LUW -> GaussDB distributed | Supporte d | Not supported | Not supported |
| To the cloud | MariaDB->MariaDB | Supporte d | Not supported | Not supported |
| To the cloud | MariaDB->MySQL | Supporte d | Supporte d | Supporte d |
| To the cloud | MariaDB->GaussDB(for MySQL) | Supporte d | Supporte d | Supporte d |
| From the cloud | MySQL->MySQL | Supporte d | Supporte d | Supporte d |
| From the cloud | MySQL->Kafka | Not supported | Not supported | Supporte d |
| From the cloud | MySQL->CSS/ES | Supporte d | Not supported | Supporte d |
| From the cloud | MySQL->MariaDB | Supporte d | Supporte d | 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 -> Kafka | Not supported | Not supported | Supporte d |

| Sync hroni zatio n Direc tion | Data Flow | Data Filtering | Additional Column | Column Processin g |
|-------------------------------|--|----------------|-------------------|--------------------|
| From the cloud | GaussDB primary/standby -> GaussDB(DWS) | Supporte d | Not supported | Not supported |
| From the cloud | GaussDB primary/standby -> GaussDB distributed | Supporte d | Not supported | Supporte d |
| From the cloud | GaussDB primary/standby -> GaussDB primary/standby | Supporte d | Not supported | Supporte d |
| From the cloud | GaussDB distributed -> MySQL | Supporte d | Not supported | Not supported |
| From the cloud | GaussDB distributed -> Oracle | Supporte d | Not supported | Supporte d |
| From the cloud | GaussDB distributed -> GaussDB(DWS) | Supporte d | Not supported | Not supported |
| From the cloud | GaussDB distributed -> Kafka | Not supported | Not supported | Supporte d |
| From the cloud | GaussDB distributed -> GaussDB distributed | Supporte d | Not supported | Supporte d |
| From the cloud | GaussDB distributed->GaussDB primary/standby | Supporte d | Not supported | Supporte d |
| 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 |

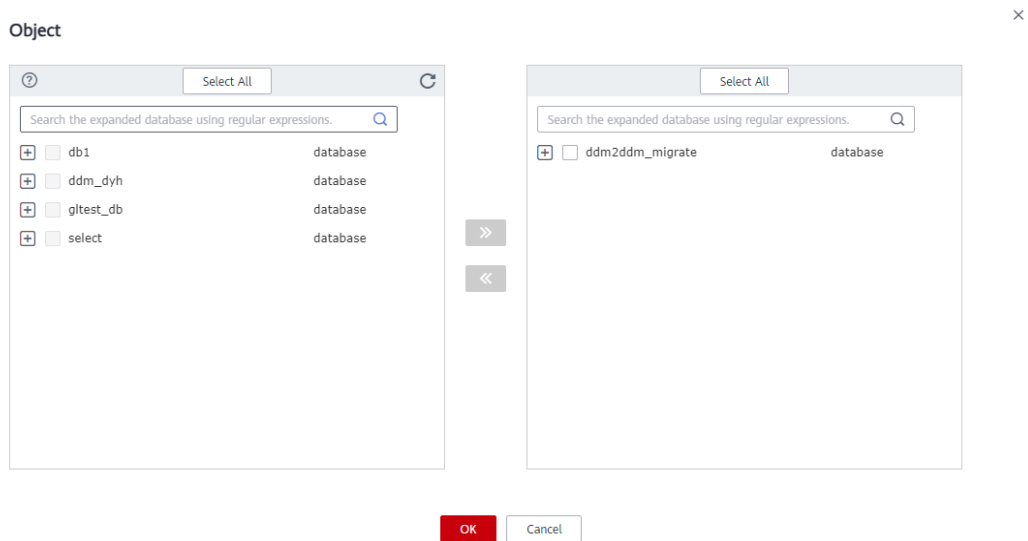
| Sync hroni zatio n Direc tion | Data Flow | Data Filtering | Additional Column | Column Processin g |
|--|--|-------------------|----------------------|--------------------------|
| From the cloud | MariaDB->MariaDB | Supporte d | Not supported | Not supported |
| Self- built -> Self- built | MySQL->Kafka | Not supported | Not supported | Supporte d |
| Self- built -> Self- built | MySQL->CSS/ES | Supporte d | Not supported | Supporte d |
| Self- built -> Self- built | Oracle -> GaussDB primary/ standby | Supporte d | Not supported | Supporte d |
| Self- built -> Self- built | Oracle -> GaussDB distributed | Supporte d | Not supported | Supporte d |
| Self- built -> Self- built | GaussDB primary/standby -> Kafka | Not supported | Not supported | Supporte d |
| Self- built -> Self- built | GaussDB distributed -> Kafka | Not supported | Not supported | Supporte d |
| Self- built -> Self- built | DB2 for LUW -> GaussDB primary/ standby | Supporte d | Not supported | Not supported |

| Sync hroni zatio n Direc tion | Data Flow | Data Filtering | Addition al Column | Column Processin g |
|--|---------------------------------------|-------------------|--------------------------|--------------------------|
| Self- built -> Self- built | DB2 for LUW -> GaussDB distributed | Supporte d | Not supported | Not supported |

Adding Synchronization Timestamp


- Step 1** On the **Additional Columns** tab of the **Process Data** page of the real-time synchronization task, click to **Select Object**.
- Step 2** In the dialog box that is displayed, select the table objects to be processed and click **OK**.

Figure 6-39 Selecting objects



- Step 3** In the **Process Data** area, enter the column name, type, and operation type to be added.

Figure 6-40 Adding information



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 4 Click **Next**.

----End

Adding Additional Columns

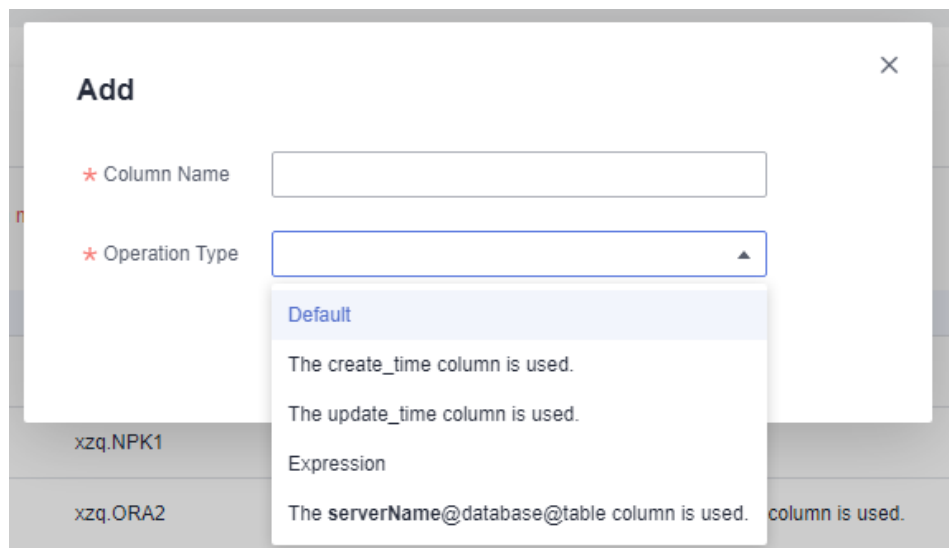
Step 1 On the **Process Data** page of the real-time synchronization task, click **Additional Columns**, locate the table to be processed, and click **Add** in the **Operation** column.

Figure 6-41 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 6-42 Operation types



 **NOTE**

- In many-to-one mapping scenarios, additional columns for data processing are required to avoid data conflicts.
- The following operation types are supported:
 - **Default:** Use the default value to fill in the new column.
 - Use the `create_time` column and `update_time` column as an example to fill the new column with the data creation time and data update time.
 - **Expression:** Use the `concat(_current_database, '@', _current_table)` expression to fill in the new column. You cannot manually enter an expression.
 - If you fill in the new column in `serverName@database@table` format, you need to enter a server name and then the database name and table name will be automatically filled in.
 - **Value:** Select a value, for example, synchronization time.
- You can apply the additional column information of the first editable table to all editable tables in batches.
- During MySQL to GaussDB(for MySQL) synchronization, if the number of columns in a single table exceeds 500, the number of additional columns added to the table may exceed the upper limit. As a result, the task fails.
- If `serverName@database@table` is used to add an additional column, this additional column will be used on the destination database as an implicit filtering condition for row comparison and value comparison by default.

Step 3 Click **Next**.

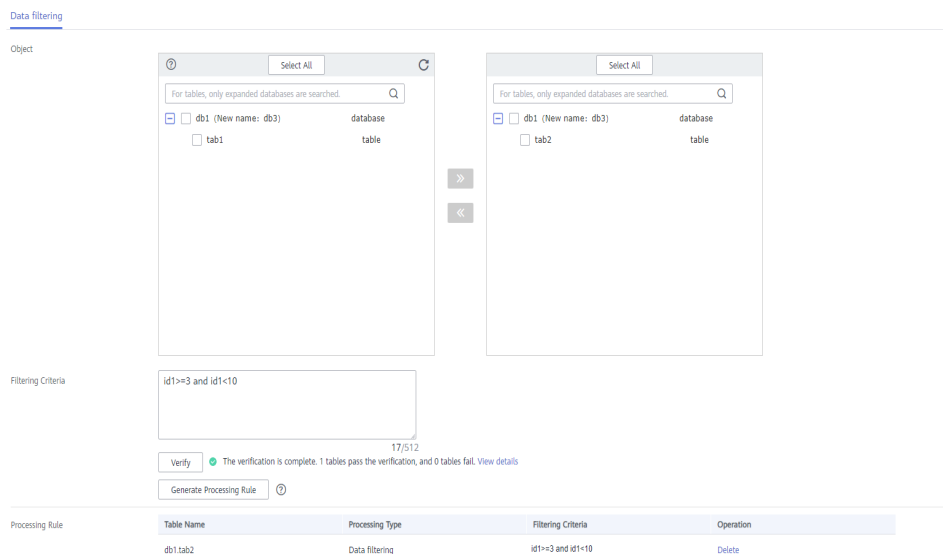
----End

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 6-43 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) and character string type must be enclosed in single quotation marks, for example, ``update` > '2022-07-13 00:00:00'` and `age >10`, ``update` ='abc'`.
- Implicit conversion rules are not supported. Enter filtering criteria of a valid data type. For example, if column `c` of an Oracle database uses characters of the `varchar2` type, the filtering criteria must be set to `c > '10'` instead of `c > 10`.
- Filter criteria cannot be configured for large objects, such as CLOB, BLOB, and BYTEA.
- Filtering rules cannot be set for objects whose database names and table names contain newline characters.
- The syntax of row-level locks, such as for update, cannot be used as filtering criteria.
- Function operations cannot be performed on column names. If function operations are performed, data may be inconsistent.
- You are not advised to set filter criteria for fields of approximate numeric types, such as FLOAT, DECIMAL, and DOUBLE.
- Do not use fields containing special characters as a filter condition.
- You are not advised to use non-idempotent expressions or functions as data processing conditions, such as SYSTIMESTAMP and SYSDATE, because the returned result may be different each time the function is called.
- The filtering rules for a synchronized table cannot be modified.

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**.

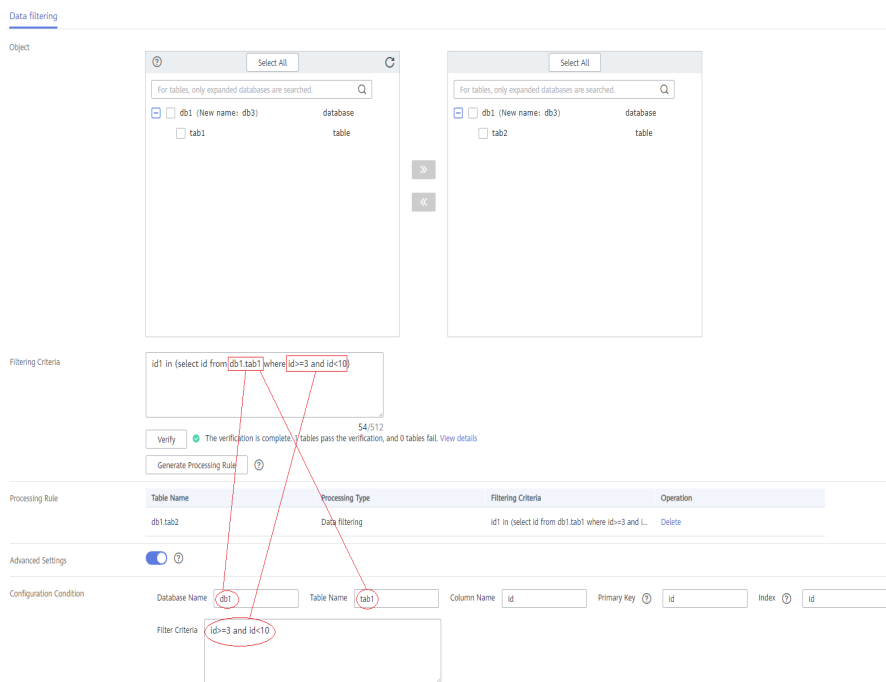
NOTE

- 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.
- Implicit conversion rules are not supported. Enter filtering criteria of a valid data type. For example, if column c of an Oracle database uses characters of the varchar2 type, the filtering criteria must be set to c > '10' instead of c > 10.
- Filter criteria cannot be configured for large objects, such as CLOB, BLOB, and BYTEA.
- Filtering rules cannot be set for objects whose database names and table names contain newline characters.
- The syntax of row-level locks, such as for update, cannot be used as filtering criteria.
- 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 6-44 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**.

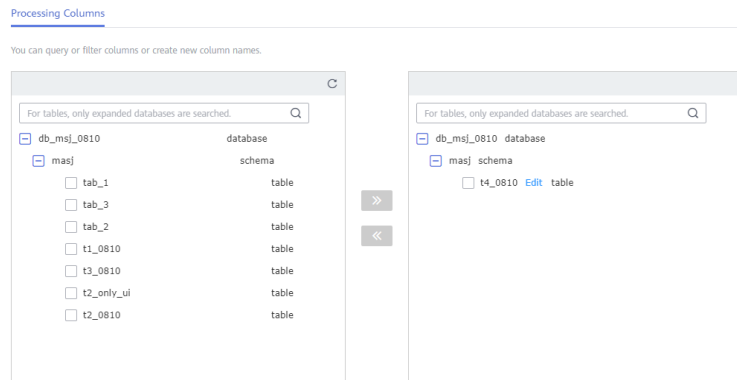
----End

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.

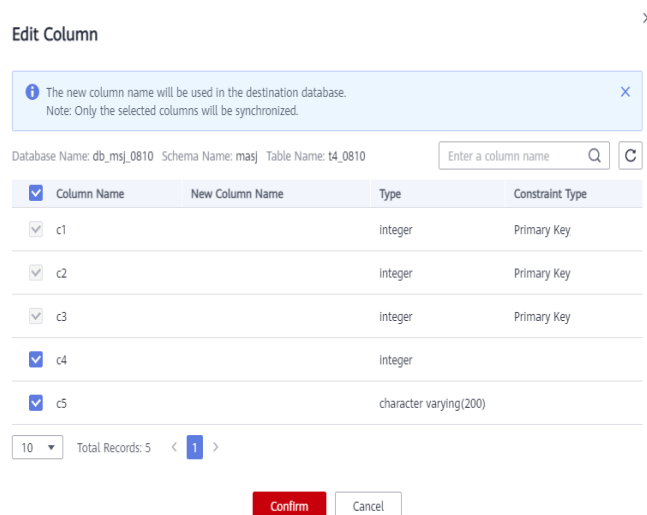
Figure 6-45 Processing Columns



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 6-46 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 selected columns are synchronized. Newly-added columns are not included in column processing.
- 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.
- In the incremental phase, DDL operations cannot be performed on filtered or mapped columns in a table.


Step 5 Click **Confirm**.

Step 6 Click **Next**.

----End

Viewing Data Filtering Results

Step 1 On the **Data Synchronization Management** page, click the task to be processed.

Step 2 Click the **Process Data** tab to view data filtering records. Click  in the upper right corner to refresh the record list.

----End

View Column Processing

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.


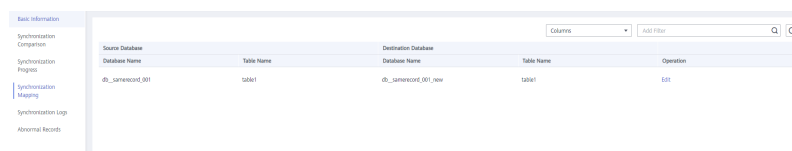
Step 2 In the navigation pane on the left, choose **Synchronization Mapping**. In the upper right corner, and select **Columns** to view column mapping records. Click  in the upper right corner to refresh the record list.

Figure 6-47 Viewing column mappings



| Source Database | | Destination Database | | Operation |
|-----------------|------------|----------------------|------------|-----------|
| Database Name | Table Name | Database Name | Table Name | |
| db_saprecore001 | table1 | db_saprecore001_new | table1 | Edit |

----End

6.6 Managing Parameters

6.6.1 Changing Task Parameters

DRS allows you to change task parameters based on service requirements.

Constraints

- You can change parameters of a synchronization task in the **Configuration, Full, Full synchronization failed, Incremental, Incremental synchronization failed, or Paused** state.
- After some parameters of a task are changed, the changes take effect only after you restart the task by referring to [Restarting a Synchronization Task](#). In addition, operations for resetting a task, resuming a task, rebuilding a task, performing primary/standby switchover, or changing the flow control in the full synchronization phase will restart task processes for the parameter changes to take effect.
- You can change parameters only for the synchronization tasks of the following engines: For details about common task parameters, see [Task Parameter Description](#).
 - MySQL->MySQL
 - MySQL->PostgreSQL
 - MySQL -> GaussDB Distributed
 - MySQL -> GaussDB Primary/Standby
 - MySQL->GaussDB(DWS)
 - MySQL->GaussDB(for MySQL)
 - MySQL->Kafka
 - MySQL->CSS/ES
 - MySQL->Oracle

Procedure

- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task.
- Step 2** Click the task name. The **Basic Information** page is displayed.
- Step 3** Choose **Parameters** on the left pane. On the **Parameters** tab page, change the parameters of the current synchronization task.

Figure 6-48 Parameters

| Parameter Name | Effective upon Restart | Value | Allowed Values | Description |
|--------------------------------|------------------------|--------|-------------------|---|
| applier_thread_num | Yes | 4 | 1-8 | Number of threads for writing data to the destination da... |
| relay_socket_timeout | Yes | 30000 | 30000-2147483647 | Timeout interval for obtaining logs from the source data... |
| shard_task_num | Yes | 2 | 1-4 | Number of sharding threads in the source database da... |
| datamove_source_socket_timeout | Yes | 120000 | 120000-2147483647 | Timeout interval for obtaining data from the source dat... |
| read_task_num | Yes | 2 | 1-4 | Number of threads for reading data from the source da... |
| increment_writer_num | Yes | 16 | 1-32 | Number of threads for writing data to the destination da... |

Step 4 Click **Save** or **Cancel**.

- To save your changes, click **Save**. In the displayed dialog box, click **Yes**.
- To cancel your changes, click **Cancel**. In the displayed dialog box, click **Yes**.

Step 5 Restart the task for the changes to take effect. You can click **Restart** in the upper right corner or restart the task by referring to [Restarting a Synchronization Task](#).

In addition, operations for resetting a task, resuming a task, rebuilding a task, performing primary/standby switchover, or changing the flow control in the full synchronization phase will restart task processes for the parameter changes to take effect.

----End

Task Parameter Description

Table 6-11 Task parameters

| Parameter | Description | Value Range | Default Value | Type | Restart Required |
|--------------------|---|---|--|------|------------------|
| applier_thread_num | Number of threads for writing data to the destination database during the full synchronization phase. | The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-4 • Small: 1-4 • Medium: 1-8 • Large: 1-18 • Ultra-large: 1-24 | The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 2 • Small: 2 • Medium: 4 • Large: 6 • Ultra-large: 12 | int | Yes |
| read_task_num | Number of threads for reading data from the source database during the full synchronization phase. | The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-2 • Small: 1-2 • Medium: 1-4 • Large: 1-8 • Ultra-large: 1-12 | The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1 • Small: 1 • Medium: 2 • Large: 4 • Ultra-large: 6 | int | Yes |

| Parameter | Description | Value Range | Default Value | Type | Restart Required |
|--------------------------------|--|---|--|------|------------------|
| shard_task_num | Number of sharding threads in the source database during the full synchronization phase. | The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-2 • Small: 1-2 • Medium: 1-4 • Large: 1-8 • Ultra-large: 1-8 | The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1 • Small: 1 • Medium: 2 • Large: 4 • Ultra-large: 4 | int | Yes |
| increment_writer_num | Number of threads for writing data to the destination database during the incremental synchronization phase. | The value range depends on the task specifications. <ul style="list-style-type: none"> • Micro: 1-8 • Small: 1-16 • Medium: 1-32 • Large: 1-128 • Ultra-large: 1-128 | The default value depends on the task specifications. <ul style="list-style-type: none"> • Micro: 4 • Small: 8 • Medium: 16 • Large: 64 • Ultra-large: 64 | int | Yes |
| relay_socket_timeout | Timeout interval for obtaining logs from the source database during the incremental synchronization phase. | 30000-2147483647 ms | 30000 ms | int | Yes |
| datamove_source_socket_timeout | Timeout interval for obtaining data from the source database during the full synchronization phase. | 120000-2147483647 ms | 120000 ms | int | Yes |

6.6.2 Viewing Parameter Change History

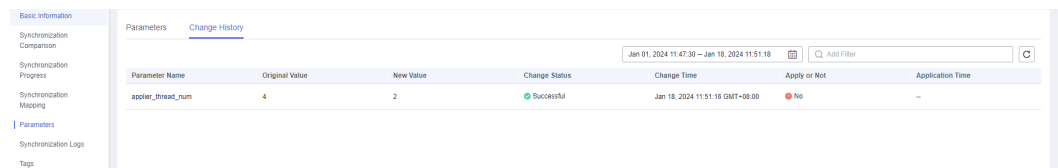
DRS allows you to view the change history of task parameters to meet service requirements.

Procedure

- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task.
- Step 2** Click the task name. The **Basic Information** page is displayed.
- Step 3** Choose **Parameters** on the left pane. On the **Change History** tab page, view the change history of parameters of the current synchronization task.

By default, the parameter change history of the last seven days can be queried.

Figure 6-49 Change History



| Parameter Name | Original Value | New Value | Change Status | Change Time | Apply or Not | Application Time |
|-------------------|----------------|-----------|---------------|---------------------------------|--------------|------------------|
| applic_thread_num | 4 | 2 | Successful | Jan 18, 2024 11:51:16 GMT+08:00 | No | -- |

----End

6.7 Task Life Cycle

6.7.1 Viewing Task Details

View the information about the synchronization tasks and synchronization instances. This section describes how to view details about a synchronization task you have created.

Prerequisites

You have logged in to the DRS console.

Procedure

NOTE

In the task list, only tasks created by the current login user are displayed. Tasks created by different users of the same tenant are not displayed.

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 On the displayed **Basic Information** page, view details about the migration task.

You can view information about the task, synchronization instance, and synchronization.

----End

6.7.2 Modifying Task Information

After a synchronization task is created, you can modify task information to identify different tasks.

The following task information can be edited:

- Task name
- Description
- SMN topic
- Synchronization delay threshold
- Number of days when an abnormal task is stopped
- Task start time

Prerequisites

You have logged in to the DRS console.

Procedure

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 On the **Basic Information** tab, locate the information to be modified in the **Task Information** area.




- You can click  to modify the task name, SMN topic, delay threshold, the time to stop abnormal tasks, and description.
 - To submit the change, click .
 - To cancel the change, click .

Table 6-12 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: !<>&'\" |

| Task Information | Description |
|---------------------------------|---|
| SMN Topic | You can apply for a topic on the SMN console and add a subscription. For details, see Simple Message Notification User Guide . |
| Synchronization Delay Threshold | The delay ranges from 0s to 3600s. NOTE If the delay threshold is set to 0, no notifications will be sent to the recipient. |
| Stop Abnormal Tasks After | The value must range from 14 to 100. The default value is 14. NOTE You can set this parameter only for pay-per-use tasks. |

- You can modify the task start time only when the task is in the **Pending start** status.
In the **Task Information** area, click **Modify** in the **Scheduled Start Time** field. On the displayed page, specify the scheduled start time and click **OK**.

Step 3 View the change result on the **Basic Information** tab.


----End

Configuring Exception Notifications

Step 1 On the **Data Synchronization Management** page, select the task to be configured.

Step 2 Click **Batch Operations** in the upper left corner and choose **Configure Exception Notification**.

Figure 6-50 Batch Operations



Step 3 In the displayed dialog box, enter the configuration information and click **Yes** to submit the configuration task.

----End

6.7.3 Modifying Connection Information

During the synchronization, you may change the password of the source or destination database. As a result, the data synchronization, data comparison, task resuming, resetting, object editing, and stopping may fail. In this case, you need to change the password on the DRS console and resume the task.

You can modify the following synchronization information:

- Database password
- Database IP address
- Database port
- Database username

Constraints

- You can change the IP address, port, and username during the incremental synchronization phase only for a DRS task with MySQL or GaussDB(for MySQL) serving as the source and IP address entered for the connection test. If the IP address, port number, or username changes due to some operations on the source database, you can use this function to update the information.
- The function of changing an IP address applies to the scenario where the IP address of the source database changes. The IP addresses before and after the change must belong to the same data instance. Otherwise, the task may fail or data may be inconsistent.
- After the connection information is changed, the change takes effect immediately, and the data in the destination database is not cleared.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click **Modify Connection Details** in the **Connection Information** area.
- Step 3** In the displayed dialog box, change the database information such as passwords of the source and destination databases and click **OK**.

----End

6.7.4 Modifying the Flow Control Mode

You can choose whether to control the flow. DRS allows you to change the flow control mode after a task is created. Currently, only the real-time synchronization scenarios listed in [Real-time Synchronization Scenarios That Support Flow Control](#) support this function.

Constraints

- The flow control mode limits the maximum traffic speed in seconds. The actual statistical value may be lower than the flow rate because the statistical value may decrease due to network fluctuation.
- The flow control mode takes effect only in the full synchronization phase.
- After the traffic rate is modified in the incremental migration phase, the modification takes effect when the task enters the full migration phase again. For example, if the traffic rate is modified and a synchronization object is added to the task, the modification takes effect in the full synchronization phase of the task.

Prerequisites

- You have logged in to the DRS console.
- A synchronization task has been created.

Method 1

Step 1 In the **Flow Control Information** area on the **Basic Information** tab, click **Modify**.

Step 2 In the displayed dialog box, modify the settings.

Figure 6-51 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 Synchronization Management** page, locate the target task and choose **More > Speed** or **Speed** in the **Operation** column.

Figure 6-52 Task List

| Task Name/ID | Status | Delay | Bill... | Dat... | DB En... | Synchr... | Created | N... | Billing ... | Description | Enter... | Operation |
|--------------|-------------|---------|---------|--------------|----------|----------------|----------------------|------|----------------|-------------|----------|------------------------|
| [Redacted] | Incremental | 300 100 | Yes | To the cl... | MySQL | Incremental | Oct 12, 2023 10:3... | VPC | Pay-per-use | | default | Edit Stop More |
| [Redacted] | Incremental | 0s | Yes | To the cl... | MySQL | Full-Increm... | Oct 11, 2023 20:3... | VPC | Pay-per-use | | default | Edit Stop More |
| [Redacted] | Incremental | 0s | Yes | To the cl... | MySQL | Full-Increm... | Oct 10, 2023 15:2... | VPC | Yearly/Monthly | | default | Edit Speed More |

Step 2 In the displayed dialog box, modify the settings.

Figure 6-53 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

Real-time Synchronization Scenarios That Support Flow Control

To the cloud

- MySQL->MySQL
- MySQL->GaussDB(for MySQL)
- MySQL -> GaussDB primary/standby
- MySQL -> GaussDB Distributed
- MySQL->PostgreSQL
- MySQL->GaussDB(DWS)
- MySQL->MariaDB
- Oracle->MySQL
- Oracle->GaussDB(for MySQL)
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- Oracle->GaussDB(DWS)
- Oracle->DDM
- Oracle->PostgreSQL
- PostgreSQL->PostgreSQL
- PostgreSQL -> GaussDB Primary/Standby
- PostgreSQL -> GaussDB Distributed
- 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
- MariaDB->MariaDB
- MariaDB->MySQL
- MariaDB->GaussDB(for MySQL)
- GaussDB(for MySQL)->GaussDB(for MySQL)

From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- MySQL->MariaDB
- GaussDB(for MySQL)->Oracle
- GaussDB(for MySQL)->Kafka
- GaussDB primary/standby -> MySQL
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> GaussDB(DWS)
- GaussDB primary/standby -> GaussDB distributed
- GaussDB primary/standby -> GaussDB primary/standby
- GaussDB distributed -> MySQL
- GaussDB distributed -> Oracle
- GaussDB distributed -> GaussDB(DWS)
- GaussDB distributed -> GaussDB distributed
- GaussDB distributed->GaussDB primary/standby
- MariaDB->MariaDB
- PostgreSQL->PostgreSQL

Self-built -> Self-built

- MySQL->CSS/ES
- MySQL->Kafka
- Oracle->Kafka
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> GaussDB primary/standby
- GaussDB distributed -> Oracle
- GaussDB distributed -> GaussDB distributed

6.7.5 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

6.7.6 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 right corner.


----End

Resuming Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be resumed.

Step 2 Click **Batch Operations** in the upper left corner and choose **Resume**.

Figure 6-54 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

6.7.7 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.

6.7.8 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. You are not advised pausing a task for more than 24 hours. For details, check the corresponding log configuration.


----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 6-55 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 Distributed
 - MySQL->GaussDB(DWS)
 - MySQL->PostgreSQL
 - MySQL->MariaDB
 - PostgreSQL->PostgreSQL
 - PostgreSQL->GaussDB(DWS)
 - PostgreSQL -> GaussDB Primary/Standby
 - PostgreSQL -> GaussDB Distributed
 - DDM->DDM
 - DDM->GaussDB(DWS)
 - Oracle->GaussDB(DWS)
 - Oracle->PostgreSQL
 - Oracle->GaussDB(for MySQL)
 - Oracle -> GaussDB primary/standby
 - Oracle -> GaussDB distributed
 - Oracle->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
 - MariaDB->MariaDB
 - MariaDB->MySQL
 - MariaDB->GaussDB(for MySQL)
 - GaussDB(for MySQL)->GaussDB(for MySQL)
- From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- MySQL->MariaDB
- DDM->Oracle
- DDM->Kafka
- DDS->MongoDB
- DDS->Kafka
- GaussDB(for MySQL)->MySQL
- GaussDB(for MySQL)->Kafka
- GaussDB(for MySQL)->CSS/ES
- GaussDB(for MySQL)->Oracle
- GaussDB(for MySQL)->GaussDB(DWS)
- GaussDB primary/standby -> MySQL
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> Kafka
- GaussDB distributed -> MySQL
- GaussDB distributed -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->PostgreSQL
- PostgreSQL->Kafka
- MariaDB->MariaDB
- Microsoft SQL Server->Kafka
- Self-built -> Self-built
 - MySQL->Kafka
 - MySQL->CSS/ES
 - MySQL -> GaussDB Primary/Standby
 - MySQL -> GaussDB Distributed
 - Oracle-> Kafka
 - Oracle -> GaussDB primary/standby
 - Oracle -> GaussDB distributed
 - GaussDB primary/standby -> Oracle
 - GaussDB primary/standby -> Kafka
 - GaussDB distributed -> Oracle
 - GaussDB distributed -> Kafka
 - PostgreSQL->Kafka
 - DB2 for LUW -> GaussDB primary/standby
 - DB2 for LUW -> GaussDB distributed
 - Microsoft SQL Server->Kafka

In addition, the following tasks can be paused during full synchronization:

- MySQL->MySQL
- MySQL->GaussDB(for MySQL)
- MySQL->GaussDB(DWS)
- MySQL->CSS/ES
- MySQL->Kafka
- GaussDB(for MySQL)->MySQL
- GaussDB(for MySQL)->GaussDB(for MySQL)
- GaussDB(for MySQL)->Kafka
- GaussDB(for MySQL)->GaussDB(DWS)
- Oracle->MySQL
- DDM->MySQL
- DDM->Oracle
- DDM->GaussDB(DWS)
- PostgreSQL->PostgreSQL
- PostgreSQL->GaussDB(DWS)
- MongoDB->DDS
- DDS->MongoDB

6.7.9 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

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 right corner.

Step 3 Perform [Step 2](#) to [Step 3](#) from method 1.

----End

Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Reset

To the cloud

- MySQL->MySQL
- MySQL->PostgreSQL
- MySQL->GaussDB(for MySQL)
- MySQL -> GaussDB primary/standby
- MySQL -> GaussDB Distributed
- MySQL->GaussDB(DWS)
- MySQL->MariaDB
- PostgreSQL->PostgreSQL
- PostgreSQL -> GaussDB Primary/Standby
- PostgreSQL -> GaussDB Distributed
- 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
- MariaDB->MariaDB
- MariaDB->MySQL
- MariaDB->GaussDB(for MySQL)
- GaussDB(for MySQL)->GaussDB(for MySQL)

From the cloud

- MySQL->MySQL
- MySQL->CSS/ES
- MySQL->Oracle
- MySQL->Kafka
- MySQL->MariaDB
- DDS->MongoDB
- DDS->Kafka
- 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 -> MySQL
- GaussDB distributed -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->PostgreSQL
- PostgreSQL->Kafka
- MariaDB->MariaDB
- Microsoft SQL Server->Kafka

Self-built -> Self-built

- MySQL->Kafka

- MySQL->CSS/ES
- MySQL -> GaussDB Primary/Standby
- MySQL -> GaussDB Distributed
- Oracle-> Kafka
- Oracle -> GaussDB primary/standby
- Oracle -> GaussDB distributed
- GaussDB primary/standby -> Oracle
- GaussDB primary/standby -> Kafka
- GaussDB distributed -> Oracle
- GaussDB distributed -> Kafka
- PostgreSQL->Kafka
- DB2 for LUW -> GaussDB primary/standby
- DB2 for LUW -> GaussDB distributed
- Microsoft SQL Server->Kafka

6.7.10 Resetting Synchronization Position

For real-time synchronization tasks from MySQL to GaussDB(DWS), DRS allows you to reset the incremental synchronization position to place the synchronization task to the time before a fault occurs and perform synchronization again.


Constraints

- Only synchronization tasks from MySQL to GaussDB(DWS) support this function.
- When a task is in the incremental phase, you need to pause the task and then reset the synchronization position.

Procedure

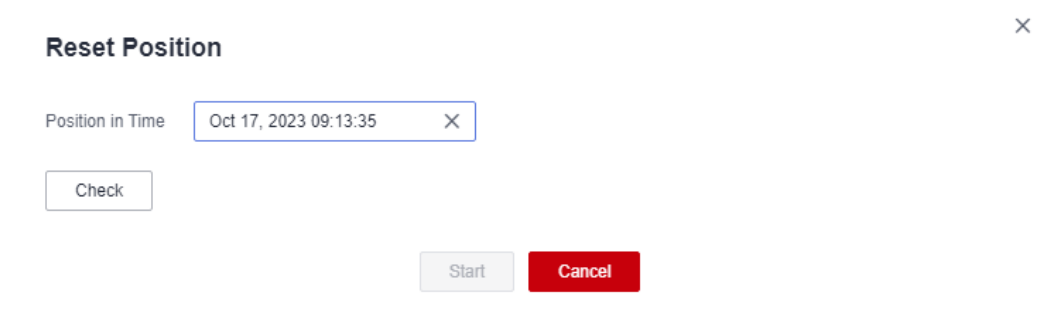
- Step 1** On the **Data Synchronization Management** page, locate the target synchronization task and choose **More > Reset Position** in the **Operation** column.

Figure 6-56 Reset Position



- Step 2** In the **Reset Position** dialog box, specify **Position in Time** and click **Check**.

Figure 6-57 Confirming dialog box



Step 3 After verifying the synchronization position, click **Start** to submit the synchronization task again.

----End

6.7.11 Restarting a Synchronization Task

DRS allows you to restart a synchronization task after task parameters are changed in [Changing Task Parameters](#).

Constraints

- Tasks in the **Full**, **Full synchronization failed**, **Incremental**, **Incremental synchronization failed**, or **Paused** state can be restarted.
- Only the synchronization tasks of the following engines can be restarted:
 - MySQL->MySQL
 - MySQL->PostgreSQL
 - MySQL -> GaussDB Distributed
 - MySQL -> GaussDB Primary/Standby
 - MySQL->GaussDB(DWS)
 - MySQL->GaussDB(for MySQL)
 - MySQL->Kafka
 - MySQL->CSS/ES
 - MySQL->Oracle

Method 1

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Restart** in the **Operation** column.

Step 2 In the displayed **Restart Task** dialog box, click **Yes**.

----End

Method 2

Step 1 On the **Data Synchronization Management** page, click the target synchronization task in the **Task Name/ID** column.

Step 2 On the **Basic Information** page, click **Restart** in the upper right corner.

Step 3 In the displayed **Restart Task** dialog box, click **Yes**.

----End

6.7.12 Skipping DDL Statements

In a synchronization task, if the destination is a GaussDB(DWS) instance, DDL operations that are performed on the source database in the incremental synchronization phase will also be performed on the GaussDB(DWS) instance. If the destination GaussDB(DWS) instance does not support the DDL statement, the synchronization task becomes faulty.


To solve this problem, you can skip the DDL statement to continue the synchronization task after rectifying the fault based on the synchronization log information.

Procedure

Step 1 Contact GaussDB(DWS) technical support to execute this statement with the same semantics in the destination database and the statement must comply with the syntax of the destination database.


Step 2 On the **Data Synchronization Management** page, locate the task you want to skip and click **Skip DDL** in the **Operation** column.

Figure 6-58 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 6-59 Skip DDL



----End

6.7.13 Performing a Switchover for a Dual-AZ Task

When you create a synchronization task, you can select the DRS task type. The DRS task can be a single-AZ task or a dual-AZ task. The dual-AZ deployment

provides HA, improving the reliability of DRS tasks. After a dual-AZ task is created, DRS creates two subtasks, each running in the primary and standby AZs. If the subtask in the primary AZ fails, DRS automatically starts the subtask in the standby AZ to continue the synchronization.

You can select the DRS task type in the following scenarios:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL->GaussDB(DWS)
 - MySQL->MariaDB
 - GaussDB(for MySQL)->GaussDB(for MySQL)
 - MariaDB->MySQL
 - MariaDB->GaussDB(for MySQL)
 - Oracle -> GaussDB Primary/Standby
 - Oracle -> GaussDB Distributed
- From the cloud
 - MySQL->Kafka
 - MySQL->MariaDB
 - GaussDB(for MySQL)->Kafka
 - GaussDB(for MySQL)->CSS/ES
 - GaussDB primary/standby -> Kafka
 - GaussDB Distributed -> Kafka
- Self-built -> Self-built
 - MySQL->Kafka
 - Oracle->Kafka
 - Oracle -> GaussDB Primary/Standby
 - Oracle -> GaussDB Distributed
 - GaussDB primary/standby -> Kafka
 - GaussDB Distributed -> Kafka

Prerequisites

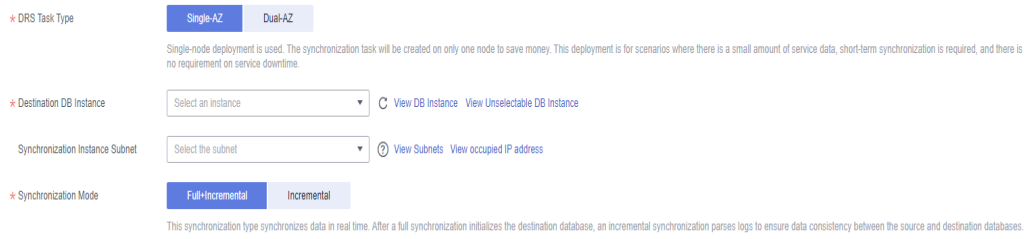
- You have logged in to the DRS console.

Scenarios

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.


Step 2 On the **Create Synchronization Instance** page, configure the task name, description, and the synchronization instance details, set **DRS Task Type** to **Dual-AZ**, and click **Next**.

Figure 6-60 Synchronization instance details




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 6-61 Primary/Standby tasks




Step 4 After a synchronization task is configured and started, DRS will start the task in the primary AZ, and the task in the standby AZ is suspended.

Figure 6-62 Before the primary/standby switchover



Step 5 If the task in the primary AZ is abnormal, DRS automatically starts the task in the standby AZ to continue the synchronization.

Figure 6-63 After the primary/standby switchover



----End

6.7.14 Cloning a Synchronization Task

DRS allows you to clone the configuration of existing synchronization tasks. However, tasks in the following status cannot be cloned:

- Creating
- Creation failed

- Configuration
- Pending start
- Starting
- Deleted

You can clone the following data flow types:

- To the cloud
 - MySQL->MySQL
 - MySQL->MariaDB
 - PostgreSQL->PostgreSQL
- From the cloud
 - MySQL->MySQL
 - MySQL->CSS/ES
 - MySQL->MariaDB
 - DDS->MongoDB
- Self-built -> Self-built
 - MySQL->CSS/ES

NOTE

- When a task is cloned, the source and destination database passwords are not cloned. You need to enter the passwords again for the new task.
- When you clone a task, the advanced settings for data filtering are not cloned. You need to set the advanced settings for the cloned task again.
- Many-to-one task cloning is not supported.
- When you clone a task that is being changed, if the change information has been saved to the database, the clone task configuration is the same as the changed configuration.
- After a clone task is created, the EIP and private IP address of the new task are different from those of the original task. You may need to configure the network to ensure that the new task can communicate with the source and destination databases.

Prerequisites

- You have logged in to the DRS console.
- A synchronization task has been created.

Procedure

Step 1 On the **Data Synchronization Management** page, select the task to be cloned and click **Clone** in the **Operation** column.

Step 2 In the displayed dialog box, confirm the new task name and click **OK**.

Step 3 After the task is submitted and the task clone is complete, the task status changes to **Configuration**. You can click **Edit** in the **Operation** column, enter the source and destination database passwords again, and edit and start the task.

----End

6.7.15 Changing Specifications

You can change the DRS task specifications based on your service requirements. After the specification change starts, the task enters the **Changing specifications** state and data synchronization is suspended. After the specification change is complete, the task is automatically resumed. 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.
- DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be upgraded for dual-AZ tasks or downgraded.
- DRS allows you to change the specifications of synchronization tasks only in the **Full**, **Full failed**, **Incremental**, or **Incremental failed** state.
- You are advised to change the task specifications during off-peak hours.
- After the specification change starts, the task is suspended. The task is automatically resumed after the change is complete.
- It takes about 5 to 10 minutes to change the task specifications.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and project.

Step 3 Choose **Databases > Data Replication Service**. The **Data Replication Service** page is displayed.

Step 4 On the **Data Synchronization Management** page, select the target task and choose **More > Change Specifications** in the **Operation** column.

Step 5 On the displayed page, select the desired specifications, perform a pre-check, and click **Next**.

Step 6 Confirm specifications.

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances, click **Change**.

To view the cost incurred by the specifications change, choose **Billing Center > Cost Bills** in the upper right corner.

- For yearly/monthly DB instances, click **Change**. On the displayed page, click **Pay**. You can change the specifications only after the payment is successful.

Step 7 View the task specification change result.

After the application is submitted, click **Back to Task List**. On the **Data Synchronization Management** page, the instance status is **Changing specifications**.

After the task status changes from **Changing specifications** to another status, you can view the instance specifications on the **Basic Information** page to check

whether the change is successful. Alternatively, you can view the change logs on the **Synchronization Logs** page to whether the change is successful.

- **change specification start:** indicates that the specification change starts.
- **change specification success:** indicates that the specifications are changed.
- **change specification failed:** indicates that the specifications fail to be changed.

Figure 6-64 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-4a98-9344-45cb22]@201 completed, total item:4, success item:4, not pass item:0 |
| 2022/12/02 19:14:32 GMT+08:00 | Info | precheck [jacecaa5-782c-4a98-9344-45cb22]@201 start |
| 2022/12/02 19:14:18 GMT+08:00 | Info | precheck [jacecaa5-782c-4a98-9344-45cb22]@201 completed, total item:4, success item:4, not pass item:0 |
| 2022/12/02 19:13:57 GMT+08:00 | Info | precheck [jacecaa5-782c-4a98-9344-45cb22]@201 start |
| 2022/12/02 19:12:45 GMT+08:00 | Info | change specification start |
| 2022/12/02 18:45:17 GMT+08:00 | Error | service LOGMANAGER failed, cause by: database log download failed, error code is '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

6.7.16 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.

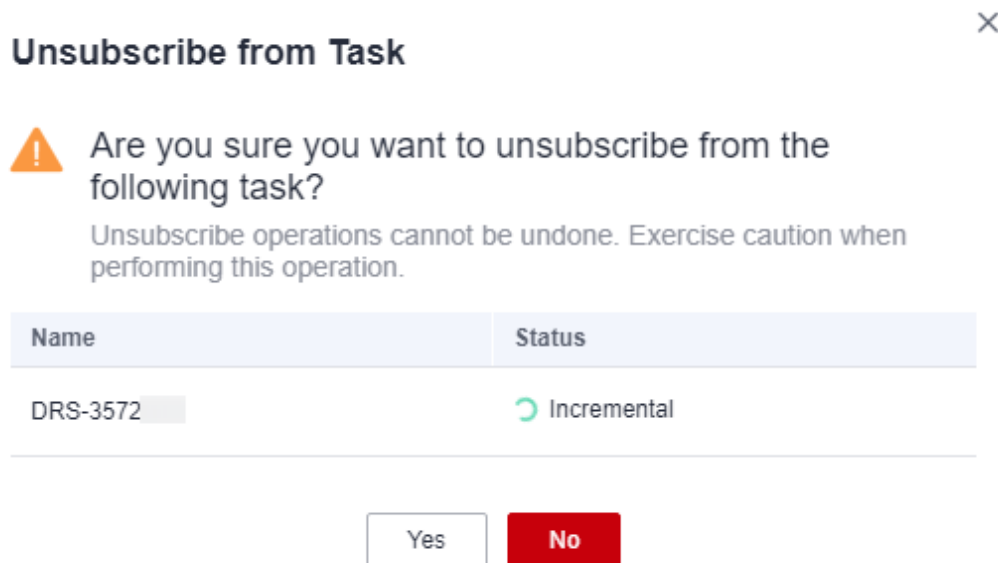
Step 2 Click  in the upper left corner and select a region and project.

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 6-65 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 In the navigation pane, choose **Orders > Unsubscriptions**.

Step 6 On the displayed page, select the order to be unsubscribed and click **Unsubscribe** in the **Operation** column.

- You can select DRS in the **Service Type** to filter all DRS orders.

Figure 6-66 Filtering all orders

The screenshot shows a table with columns: NameID, Service Type, Current Configuration, Region, Enterprise Project, and Effective Time. A search filter is applied to the Service Type column, showing a dropdown menu with options: (Select all), Elastic Cloud Server, Cloud Backup and Recovery, Relational Database Service, Data Replication Service (selected), Virtual Private Cloud, GaussDB for MySQL, and DevCloud. The table lists several orders, including Data Replication, Data Replication Service, and GaussDB for MySQL.

| NameID | Service Type | Current Configuration | Region | Enterprise Project | Effective Time |
|--------|--------------------------|------------------------------|--------|--------------------|---------------------------------|
| | Data Replication | | | default | Dec 14, 2022 15:17:38 GMT+08:00 |
| | Data Replication | | | default | Dec 14, 2022 11:59:07 GMT+08:00 |
| | Data Replication | | | default | Dec 13, 2022 16:58:21 GMT+08:00 |
| | Data Replication | | | default | Dec 13, 2022 18:00:27 GMT+08:00 |
| | Data Replication | | | default | Dec 13, 2022 17:10:05 GMT+08:00 |
| | Data Replication Serv... | DRS Sync VM (y2) | | default | Dec 13, 2022 16:05:48 GMT+08:00 |
| | GaussDB for MySQL | GaussDB (for MySQL) Cluster | | default | Dec 12, 2022 22:55:38 GMT+08:00 |
| | Cloud Backup and Re... | Cloud Server Backup Vaul2... | | default | Dec 12, 2022 15:56:35 GMT+08:00 |

- Alternatively, search for target orders by name, order No., or ID in the search box.

Step 7 On the displayed page, confirm the order to be unsubscribed from and select a reason. Then, click **Confirm**.

For unsubscription details, see [Unsubscription Rules](#).

Step 8 In the displayed dialog box, click **Yes**.

NOTE

After a DRS instance is unsubscribed, the DRS task ends immediately. Ensure that data synchronization is complete or the DRS instance is no longer used.

----End

6.7.17 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 6-67 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

6.7.18 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 6-68 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

6.7.19 Importing Tasks in Batches

DRS allows you to import an Excel file to create synchronization tasks in batches. You need to fill in a template with basic task information, source and destination database information, and objects to be synchronized, and then import the template.

Constraints

- Only real-time synchronization tasks from MySQL to MySQL, from MySQL to GaussDB Primary/Standby, from MySQL to GaussDB Distributed, from Oracle to GaussDB Primary/Standby, from Oracle to GaussDB Distributed, from GaussDB Primary/Standby to GaussDB Primary/Standby, from GaussDB Primary/Standby to GaussDB Distributed, from GaussDB Primary/Standby to Kafka, from GaussDB Distributed to GaussDB Distributed, from GaussDB Distributed to GaussDB Primary/Standby, and from GaussDB Distributed to Kafka can be imported.

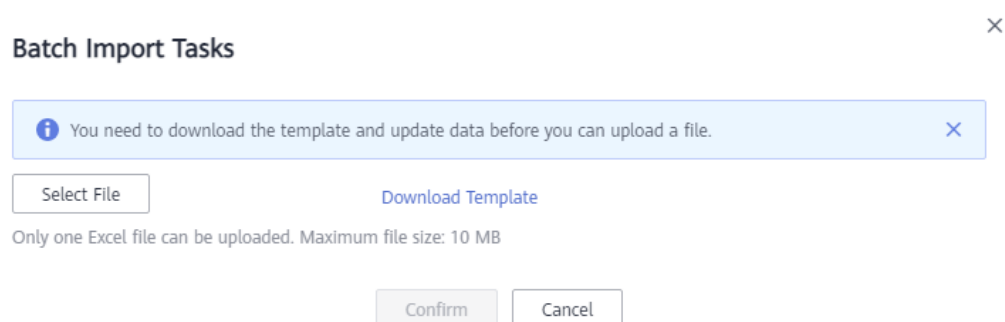
- When editing the template for importing tasks, clear the example content and enter the required values.
- When specifying the source and destination databases, you can set only mandatory parameters. You need to ensure that the subnets of the source and destination databases can communicate with each other.
- The import template is designed based on the API and GUI information. You need to set the parameters based on the description and ensure the parameter accuracy.
- After the import template is uploaded, the system applies for creating DRS instances. Some parameters cannot be modified. Exercise caution when performing this operation.
- After the import template is uploaded, the system automatically saves all parameters for the task and performs steps such as connection test, object selection, and pre-check. After the pre-check is successful, the task is in the **Configuration** state. You can click **Start Now** to start the task, or click **Edit** to modify other parameters.
- You cannot schedule the start time of a task in the import template. After the pre-check for the task is successful, you can click **Edit** on the GUI to access the startup page and set a scheduled start time.
- The system cannot verify the correctness of all parameters in the import template. You are advised to click **Edit** to enter the page, confirm the correctness of the data saved in each step, and then start the task.
- The import template is a common template. One sheet is used to create one task. To create tasks in batches, copy multiple sheets and set parameters.

Procedure

Step 1 On the **Data Synchronization Management** page, click **Batch Import Tasks** in the upper left corner.

Step 2 On the **Batch Import Tasks** dialog box, click **Download Template**.

Figure 6-69 Batch Import Tasks

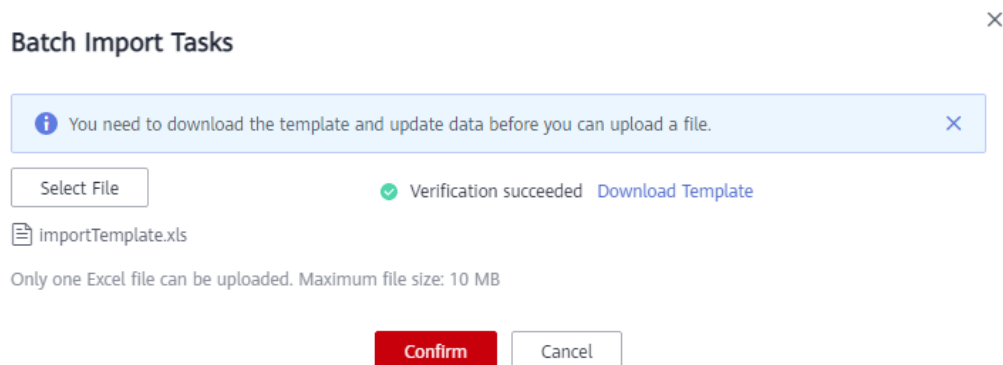


Step 3 Edit the downloaded Excel template and enter task information as required.

Step 4 Click **Select File** to upload the template.

Step 5 After the upload and verification are successful, click **Confirm** to create tasks in batches.

Figure 6-70 Import succeeded



NOTE

- DRS does not save the uploaded Excel file and only parses the content in the file.
- The IP addresses, domain names, usernames, and passwords of the source and destination databases are encrypted and stored in the system until a task is deleted.

----End

6.7.20 Task Statuses

Synchronization statuses indicate different synchronization phases.

Table 6-13 lists synchronization task statuses and descriptions.

Table 6-13 Task status description

| Status | Description |
|----------------------|--|
| Creating | A synchronization instance is being created. |
| Task creation failed | Failed to create a real-time synchronization instance. |
| Configuration | The synchronization instance is successfully created, but the synchronization task is not started. You can continue to configure the task. |
| Frozen | Instances are frozen when the account balance is less than or equal to \$0. |
| Pending start | The scheduled synchronization task has been delivered to the synchronization instance, waiting for the synchronization instance to start the synchronization task. |
| Starting | The task is being started. |
| Start failed | A real-time synchronization task fails to be started. |
| Full synchronization | A full synchronization task is being performed. |

| Status | Description |
|------------------------------------|--|
| Full synchronization failed | A full synchronization task fails. |
| Incremental synchronization | An incremental synchronization task is being performed. |
| Incremental synchronization failed | An incremental synchronization task fails. |
| Modifying task | The synchronization object is being modified. |
| Modifying task failed | The synchronization object fails to be modified. |
| Fault rectification | A synchronization instance is faulty and the system automatically restores the synchronization task. |
| Paused | The real-time synchronization task has been paused. |
| Cloning | A synchronization task is being cloned. |
| Cloning failed. | The clone synchronization task fails. |
| Changing specifications | The instance specifications are being changed. After the change is complete, the task is automatically restored. |
| Specification change failed | The instance specifications fail to be changed. |
| AZ switchover failed | The AZ switchover fails for the real-time synchronization task. |
| Task stopping | The synchronization instance and resources used for executing the synchronization task are being released. |
| Completing | A synchronization instance and resources are being released. |
| Stopping task failed | The synchronization instance and resources fail to be released. |
| Completed | The task is completed and the synchronization instance is released. |

 **NOTE**

- If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you reconfigure the task, DRS applies for resources again.
- Deleted synchronization tasks are not displayed in the status list.

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 20 tags.

Adding a Tag

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** In the navigation pane on the left, choose **Tags**.
- Step 3** On the **Tags** tab, click **Add/Edit Tags**. In the displayed dialog box, enter a tag key and value, click **Add**, and click **OK**.

Add/Edit Tags



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.

You can add 19 tags more tags.

- 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 must be unique. It must consist of 1 to 128 characters and can include letters, digits, spaces, and the following characters: `._:=-+@`. It cannot start or end with a space, or start with `_sys_`.
- The tag value can be empty. It cannot start or end with a space and can contain 0 to 255 characters, including letters, digits, spaces, and special characters `._:/=-+@`.

Step 4 View and manage the tag on the **Tags** page.

----End

Editing a Tag

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 In the navigation pane on the left, choose **Tags**.

Step 3 On the **Tags** page, click **Add/Edit Tags**. In the displayed dialog box, modify the tag and click **OK**.

----End

Delete a Tag

Step 1 On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.

Step 2 In the navigation pane on the left, choose **Tags**.

Step 3 On the **Tags** page, locate the tag to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

Step 4 After the tag is deleted, it will no longer be displayed on the **Tags** page.
----End

8 Connection Diagnosis

If a DRS instance fails to be connected to the source or destination database during connection testing, DRS provides the quick diagnosis function and returns the diagnosis result.

- You can perform connection diagnosis only on the task node whose database information is obtained by entering an IP address or selecting a task node on the GUI. DN diagnosis of GaussDB is not supported.
- In cluster or multi-AZ task scenarios, diagnosis can be performed only on the node of the primary task.

Prerequisites

- You have logged in to the DRS console.
- A task has been created.

Procedure

Step 1 On the task management page, click the target task name in the **Task Name/ID** column.

Step 2 On the **Configure Source and Destination Databases** page, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the DRS instance.

If the connection testing fails, click **Quick Diagnosis** on the right of the failure information to diagnose the fault.

Figure 8-1 Quick Diagnosis

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Database Username

Database Password

SSL Connection

● The network connection between the replication instance and database is faulty. [View details](#)

Step 3 View the diagnosis result on the displayed **Diagnosis Details** dialog box. The result includes the packet loss rate and port check result.

Figure 8-2 Diagnosis Details

| Diagnosis Details × | | |
|--|----------------------|---|
| IP Address or Domain Name | Packet Loss Rate (%) | Port Check |
| <input type="text"/> | 100 | ● Failed |

----End

9 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

10 Interconnecting with CTS

10.1 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 10-1 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 |

10.2 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

11 Interconnecting with Cloud Eye

11.1 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 11-1](#) lists the DRS performance metrics.

Table 11-1 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 apply_ddls_rate NOTE apply_ddls_rate is replaced by apply_ddls_number after December 2022. | DDL TPS | Total number of DDL transactions written into the destination database. | ≥ 0 transaction | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |

| Metric ID | Metric Name | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|----------------------------|------------------------------------|---|---|--|--------------------------------|
| apply_latency | Replication Delay | Delay (in milliseconds) of data replay. | ≥ 0 ms | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |
| apply_average_execute_time | Average Transaction Execution Time | Average execution time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is millisecond. | ≥ 0 ms | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |
| apply_average_commit_time | Average Transaction Commit Time | Average commit time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is ms. | ≥ 0 ms | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |
| apply_current_state | Synchronization Status | This metric is the synchronization status of the current kernel data (10: abnormal; 1: idle; 2: DML; 3: DDL), instead of the task status. | 10: abnormal 1: idle 2: DML is executed. 3: DDL is executed. | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |

| Metric ID | Metric Name | Description | Value Range | Monitored Object | Monitoring Interval (Raw Data) |
|----------------------|-------------------------|---|---------------------------------------|--|--------------------------------|
| apply_thread_workers | Synchronization Threads | Number of working threads for data synchronization | ≥ 0 | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |
| apply_job_status | Task Status | Status of the current task. (0: normal; 1: abnormal; 2: paused) | 0: normal 1: abnormal 2: paused | Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances | 1 minute |

Dimensions

| Key | Value |
|-------------|-----------------|
| instance_id | DRS instance ID |

11.2 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 11-1 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 11-2 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

11.3 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.

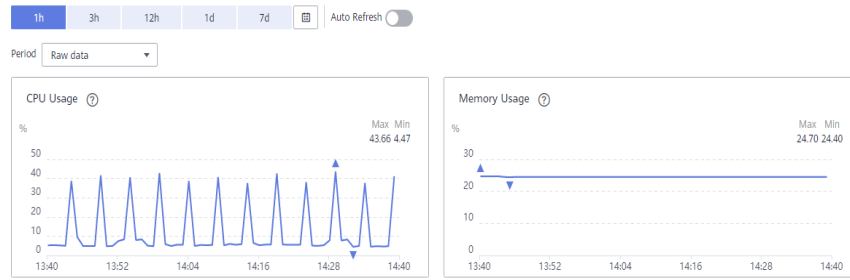
Step 5 On the displayed page, click **View Metric** in the upper right corner of the page to go to the Cloud Eye console.

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 11-3 Viewing monitoring metrics



----End

12 Interconnecting with LTS

12.1 Log Reporting

Scenarios

If you enable log reporting, all logs generated by DRS instances (including real-time migration, backup migration, real-time synchronization, real-time disaster recovery, and workload replay instances) are uploaded to Log Tank Service (LTS) for management.

Precautions


- After this function is enabled, all logs of the task are reported by default.
- This request does not take effect immediately. There is a delay of about 10 minutes.
- You will be billed for this function. For details, see [LTS Pricing Details](#).
- Ensure that there are available LTS log groups and log streams in the same region as your instance.

For more information about log groups and log streams, see [Log Management](#).

- After this function is disabled, you will not be billed anymore.

Enabling or Disabling Log Reporting

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and project.

Step 3 Choose **Database > Data Replication Service**. The **Data Replication Service** page is displayed.

Step 4 Take real-time migration as an example. On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 5 On the **Basic Information** page, click **Migration Logs** on the left.

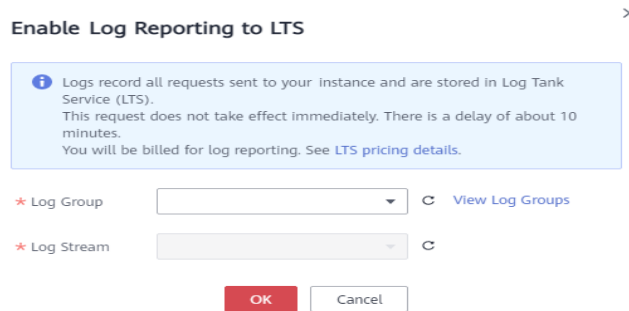
Step 6 Click  next to **Report Logs to LTS** in the upper part of the page.

Step 7 Select an LTS log group and log stream and click **OK**.

 **NOTE**

This request does not take effect immediately. There is a delay of about 10 minutes.

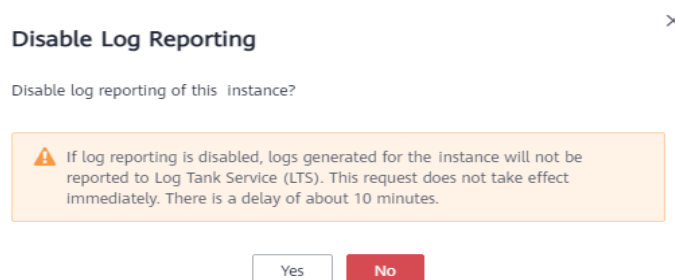
Figure 12-1 Enabling audit log reporting to LTS



Step 8 To disable or modify log reporting, click the toggle switch next to **Report Logs to LTS** or click **Edit** next to the **Report Logs to LTS** toggle switch.

- Modifying log reporting: Click **Edit** next to the **Report Logs to LTS** toggle switch. In the displayed dialog box, select the LTS log group and log stream again and click **OK**.
- Disabling log reporting: Click the toggle switch next to **Report Logs to LTS**. In the displayed dialog box, click **OK**.

Figure 12-2 Disabling log reporting to LTS



----End

12.2 Viewing and Downloading Logs

Scenarios

If you have enabled log reporting to LTS for a DRS task in [Log Reporting](#), you can analyze logs, search for logs, visualize logs, download logs, and view real-time logs on the LTS console.

Viewing Logs Reported to LTS


- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Under **Management & Governance**, click **Log Tank Service**.
- Step 4** In the **Log Groups** area, locate a target log group and click its name. For details about LTS, see [Log Tank Service User Guide](#).

Figure 12-3 Viewing log details




Table 12-1 Log field description

| Name | Type | Description |
|----------------|--------|---|
| _resource_id | String | Resource ID. The value is fixed to projectId for DRS. |
| _resource_name | String | Resource name. The value is fixed to DRS . |
| _service_type | String | Service type. The value is fixed to Data Replication Service . |

----End

Downloading Logs Reported to LTS


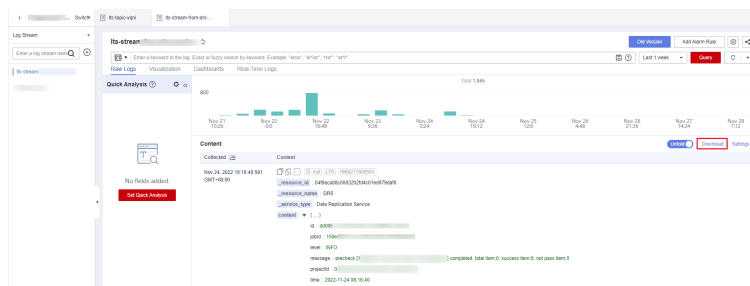
- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Under **Management & Governance**, click **Log Tank Service**.
- Step 4** In the **Log Groups** area, locate a target log group and click its name.
- Step 5** Click **Download** on the right to download logs. For details about LTS, see [Log Tank Service User Guide](#).

Figure 12-4 Downloading logs



-----End

13 Operation Reference in Synchronization Scenarios

13.1 Kafka Message Format

Data synchronized to the Kafka cluster is stored in Avro, JSON, and JSON-C formats. For details about the data formats supported by different data flow scenarios, see [Table 13-1](#).

Table 13-1 Data formats

| Data Flow | Avro | JSON | JSON-C |
|----------------------------------|---------------|-----------|---------------|
| MySQL -> Kafka | Supported | Supported | Supported |
| Oracle -> Kafka | Supported | Supported | Not supported |
| DDS->Kafka | Not supported | Yes | Not supported |
| PostgreSQL -> Kafka | Supported | Supported | Not supported |
| GaussDB(for MySQL) -> Kafka | Supported | Supported | Supported |
| GaussDB primary/standby -> Kafka | Supported | Supported | Not supported |
| GaussDB distributed -> Kafka | Supported | Supported | Not supported |
| Microsoft SQL Server->Kafka | Supported | Supported | Not supported |

Avro

For details about the schema definition in Avro format, see [record.rar](#). After data is synchronized to the Kafka cluster, parse data based on the definition of the Avro schema.

JSON

For details about the JSON format from MySQL and GaussDB (MySQL) to Kafka, see [Table 13-2](#). For details about the JSON format from DDS to Kafka, see [Table 13-3](#). For details about the JSON format from PostgreSQL, GaussDB, Microsoft SQL Server and Oracle to Kafka, see [Table 13-4](#).

Table 13-2 Parameters for synchronizing from MySQL to Kafka

| Parameter | Description |
|-----------|--|
| mysqlType | Field name and type in the source table. |
| id | Sequence number of an event operation defined in DRS. The value increases monotonically. |
| es | The time when the record is generated in the source database. The value is a 13-digit Unix timestamp in milliseconds. |
| ts | The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds. |
| database | Database name |
| table | Table name. |
| type | Operation type, for example, DELETE, UPDATE, INSERT, and DDL. For full synchronization, the value can be INIT or INIT_DDL. |
| isDdl | Whether the operation is a DDL operation. |
| sql | A DDL-defined SQL statement. The value is "". |
| sqlType | JDBC type of the fields in the source table. |
| data | The latest data, which is a JSON array. If the value of type is INSERT , the latest data is inserted. If the value of type is UPDATE , the latest data is updated. |
| old | Old data. If the value of type is UPDATE , the data is old. If the value of type is DELETE , the data is deleted. |
| pkNames | Primary key name |

```
{
  "mysqlType":{
    "c11":"binary",
    "c10":"varchar",
    "c13":"text",
```

```

    "c12": "varbinary",
    "c14": "blob",
    "c1": "varchar",
    "c2": "varbinary",
    "c3": "int",
    "c4": "datetime",
    "c5": "timestamp",
    "c6": "char",
    "c7": "float",
    "c8": "double",
    "c9": "decimal",
    "id": "int"
  },
  "id": 27677,
  "es": 1624614713000,
  "ts": 1625058726990,
  "database": "test01",
  "table": "test ",
  "type": "UPDATE",
  "isDdl": false,
  "sql": "",
  "sqlType": {
    "c11": -2,
    "c10": 12,
    "c13": -1,
    "c12": -3,
    "c14": 2004,
    "c1": 12,
    "c2": -3,
    "c3": 4,
    "c4": 94,
    "c5": 93,
    "c6": 1,
    "c7": 6,
    "c8": 8,
    "c9": 3,
    "id": 4
  },
  "data": [
    {
      "c11": "",
      "c10": "Huawei Cloud huaweicloud",
      "c13": "asfiajhfi939-0239uoituqorjqoirfojdjfrnioweioiwqjroqwjrowqjoiqgoiegnkjgoi23roiugou
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 13-3 Parameters for synchronizing data from DDS to Kafka

| Parameter | Description |
|-----------|--|
| id | Sequence number of an event operation defined in DRS. The value increases monotonically. |
| op | Operation type, such as DELETE, UPDATE, INSERT, and DDL. |
| dbType | Source database type: MongoDB |
| db | Database name. |
| coll | Collection name. |
| value | Change value of a record. |
| where | Change condition of a record. |

| Parameter | Description |
|-------------|--|
| recordType | Record type, such as insert, update, replace, and doc. update and replace indicate the UPDATE operation in op . doc indicates that the DELETE operation in op deletes document data instead of view data. |
| extra | Extended field. The value is the same as that of recordType and this parameter is used as an extended oplog record. |
| es | Commit time of a record. The value is a 13-digit Unix timestamp, in milliseconds. |
| ts | The time when the data is written to the destination Kafka. The value is a 13-digit Unix timestamp in milliseconds. |
| clusterTime | Timestamp of the oplog entry associated with the event. The value is in the format of timestamp:incr. timestamp is the Unix timestamp (unit: second), and incr is the command execution sequence in a second. |

```
// insert operation
{
  "id": 256,
  "op": "INSERT",
  "dbType": "MongoDB",
  "db": "ljk",
  "coll": "ljk",
  "value": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\"), \"c1\": \"baz\", \"tags\": [\"mongodb\", \"database\", \"NoSQL\"]}",
  "where": null,
  "recordType": "insert",
  "extra": "insert",
  "es": 1684315111439,
  "ts": 1684315111576,
  "clusterTime": "1684344064:1"
}

// replace operation
{
  "id": 340,
  "op": "UPDATE",
  "dbType": "MongoDB",
  "db": "ljk",
  "coll": "ljk",
  "value": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\"), \"c1\": \"sss\"}",
  "where": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
  "recordType": "replace",
  "extra": "replace",
  "es": 1684315951831,
  "ts": 1684315951961,
  "clusterTime": "1684344904:9"
}

// value update operation
{
  "id": 386,
  "op": "UPDATE",
```

```

"dbType": "MongoDB",
"db": "ljx",
"coll": "ljx",
"value": "{\"$set\": {\"c1\": \"aaa\"}}",
"where": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
"recordType": "update",
"extra": "update",
"es": 1684316412008,
"ts": 1684316412146,
"clusterTime": "1684345365:1"
}

// key update operation
{
  "id": 414,
  "op": "UPDATE",
  "dbType": "MongoDB",
  "db": "ljx",
  "coll": "ljx",
  "value": "{\"$unset\": {\"c1\": true}, \"$set\": {\"column1\": \"aaa\"}}",
  "where": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
  "recordType": "update",
  "extra": "update",
  "es": 1684316692054,
  "ts": 1684316692184,
  "clusterTime": "1684345648:1"
}

// remove operation
{
  "id": 471,
  "op": "DELETE",
  "dbType": "MongoDB",
  "db": "ljx",
  "coll": "ljx",
  "value": "{\"_id\": ObjectId(\"64650cf67dc36a464e76e583\")}",
  "where": null,
  "recordType": "doc",
  "extra": "doc",
  "es": 1684317252747,
  "ts": 1684317252869,
  "clusterTime": "1684346209:1"
}

```

Table 13-4 Parameters for synchronizing from other databases to Kafka

| Parameter | Description |
|------------|---|
| columnType | Field name and data type in the source table NOTE <ul style="list-style-type: none"> The data type does not contain the length and precision. This parameter is left blank when dbType is set to Oracle or Microsoft SQL Server. |
| dbType | Source database type |
| schema | Schema name. |
| opType | Operation type, such as DELETE, UPDATE, INSERT, and DDL. |

| Parameter | Description |
|-----------|---|
| id | Sequence number of an event operation defined in DRS. The value increases monotonically. |
| es | The source DB engine types are as follows: GaussDB primary/standby: commit time of the 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. Microsoft SQL Server: commit time of a record. The value is a 13-digit Unix timestamp, in milliseconds. |
| ts | The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds. |
| database | Database name. This parameter is left blank when dbType is set to Oracle . |
| table | Table name. |
| type | Operation type, such as DELETE, UPDATE, INSERT, and DDL. |
| isDdl | Whether the operation is a DDL operation. |
| sql | A DDL-defined SQL statement. The value is "". |
| sqlType | JDBC type of the fields in the source table. |
| data | The latest data, which is a JSON array. If the value of type is INSERT , the latest data is inserted. If the value of type is UPDATE , the latest data is updated. |
| old | Old data. If the value of type is UPDATE , the data is old. If the value of type is DELETE , the data is deleted. |
| pkNames | Primary key name |

```
{
  "columnType": {
    "timestamp_column": "timestamp without time zone",
    "tstzrange_column": "tstzrange",
    "int4range_column": "int4range",
    "char_column": "character",
    "jsonb_column": "json",
    "boolean_column": "boolean",
    "bit_column": "bit",
    "smallint_column": "smallint",
    "bytea_column": "bytea"
  },
}
```

```
"dbType": "GaussDB Primary/Standby",
"schema": "schema01",
"opType": "UPDATE",
"id": 332,
"es": 1639626187000,
"ts": 1639629261915,
"database": "database01",
"table": "table01",
"type": "UPDATE",
"isDdl": false,
"sql": "",
"sqlType": {
  "timestamp_column": 16,
  "tstzrange_column": 46,
  "int4range_column": 42,
  "char_column": 9,
  "jsonb_column": 22,
  "boolean_column": 8,
  "bit_column": 20,
  "smallint_column": 2,
  "bytea_column": 15
},
"data": [
  {
    "timestamp_column": "2021-12-16 12:31:49.344365",
    "tstzrange_column": "(\"2010-01-01 14:30:00+08\", \"2010-01-01 15:30:00+08\")",
    "int4range_column": "[11,20)",
    "char_column": "g",
    "jsonb_column": "{\"key1\": \"value1\", \"key2\": \"value2\"}",
    "boolean_column": "false",
    "bit_column": "1",
    "smallint_column": "12",
    "bytea_column": "62797465615f64617461"
  }
],
"old": [
  {
    "timestamp_column": "2014-07-02 06:14:00.742",
    "tstzrange_column": "(\"2010-01-01 14:30:00+08\", \"2010-01-01 15:30:00+08\")",
    "int4range_column": "[11,20)",
    "char_column": "g",
    "jsonb_column": "{\"key1\": \"value1\", \"key2\": \"value2\"}",
    "boolean_column": "true",
    "bit_column": "1",
    "smallint_column": "12",
    "bytea_column": "62797465615f64617461"
  }
],
"pkNames": null
}
```

JSON-C

JSON-C is similar to JSON. The difference lies in the delete operation. JSON data is stored in old, and JSON-C is stored in data. Data of the timestamp type is converted into a character string in the format of yyyy-mm-dd hh:mm:ss.

For details, see [Table 13-5](#).

Table 13-5 JSON-C parameter description

| Parameter | Description |
|-----------|--|
| mysqlType | Field name and type in the source table. |
| id | Sequence number of an event operation defined in DRS. The value increases monotonically. |
| es | The time when the record is generated in the source database. The value is a 13-digit Unix timestamp in milliseconds. |
| ts | The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds. |
| database | Database name. For the Oracle database, set this parameter to schema . |
| table | Table name. |
| type | Operation type, such as DELETE, UPDATE, INSERT, and DDL. |
| isDdl | Whether the operation is a DDL operation. |
| sql | A DDL-defined SQL statement. The value is "". |
| sqlType | JDBC type of the fields in the source table. |
| data | Latest data, which is a JSON array. If type is set to INSERT , this parameter indicates the latest inserted data. If type is set to UPDATE , this parameter indicates the latest updated data. If type is set to DELETE , this parameter indicates the deleted data. |
| old | Old data. If type is set to UPDATE , the value indicates the data before update. If type is set to INSERT , the value is null . |
| pkNames | Primary key name |

Common Escape Characters in JSON

Table 13-6 Escape Character

| Character | Escape character |
|-----------|------------------|
| < | \u003c |
| = | \u003d |
| > | \u003e |
| & | \u0026 |
| ' | \u0027 |

13.2 Kafka Authentication

PLAINTEXT

No security authentication mode is available. You only need to enter the IP address and port for connection.

Figure 13-1 PLAINTEXT

Destination Database

IP Address or Domain Name ?

Ensure that the entered addresses belong to the same DB instance.

Security Protocol ?

SASL_PLAINTEXT

The SASL mechanism is used to connect to Kafka, and you need to configure SASL parameters.

Figure 13-2 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 13-7 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 13-3 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 13-8 Parameter settings

| Parameter | Description |
|-----------------------------------|--|
| Truststore Certificate | SSL certificate with the file name extension .jks. For details about how to generate a certificate file, see SSL Certificate . |
| Truststore Certificate Password | Password of the certificate |
| 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 13-4 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 | <input type="text" value="SASL_SSL"/> | ? |
| SASL Mechanisms | <input type="text" value="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"/> | | |

13.3 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.

The naming rule of a replication slot depends on the database kernel version (you can run the **select working_version_num();** command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is replication_slot_unique_ID. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is drs_unique_ID. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the **task node id is ***** log on the [Synchronization Logs](#) page.

Prerequisites

Common users do not have the permission to perform the execute direct operation. To delete streaming replication slots, contact GaussDB O&M personnel.

Procedure

Step 1 Log in to each primary DN node of the GaussDB distributed instance as the user used when you tested the connectivity between the DRS instance and the GaussDB distributed instance.

Step 2 Run the following statement to query the streaming replication slot name of the database object selected for the synchronization task:

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

Step 3 Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

Step 4 Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where database = 'database';
```

If the query result is empty, the streaming replication slot is deleted.

Step 5 Repeat the preceding operations to ensure that the streaming replication slot on each primary DN is deleted.

----End

13.4 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.

The naming rule of a replication slot depends on the database kernel version (you can run the **select working_version_num();** command to view the version). If the version is earlier than 92556, serial decoding is used by default. The naming rule of a replication slot is `replication_slot_unique_ID`. If the version is 92556 or later, parallel decoding is used by default. The naming rule of a replication slot is `drs_unique_ID`. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the **task node id is ***** log on the [Synchronization Logs](#) page.

Prerequisites

Common users do not have the permission to perform the execute direct operation. To delete streaming replication slots, contact GaussDB O&M personnel.

Procedure

Step 1 Log in to the GaussDB primary/standby instance as the user used when you tested connectivity between the DRS instance and the GaussDB primary/standby instance.

Step 2 Run the following statement to query the streaming replication slot name of the database object selected for the synchronization task:

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

Step 3 Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

Step 4 Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where database = 'database';
```

If the query result is empty, the streaming replication slot is deleted.

----End

13.5 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.

The naming rule of a replication slot is *drs_unique_ID*. To obtain the unique ID, replace the hyphen (-) in the task node ID with an underscore (_). You can find the node ID in the **task node id is ***** log on the [Synchronization Logs](#) page.

Clearing the Logical Replication Slot of the Source Database

Step 1 Log in to the source database as the source database user used in the synchronization task.

- Step 2** Query the name of the streaming replication slot of the database object selected in the synchronization task.

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

- Step 3** Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

- Step 4** Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where slot_name = 'slot_name';
```

If the query result is empty, the streaming replication slot is deleted.

----End

Synchronizing Sequence Values

If sequence objects are not synchronized or the destination database is GaussDB, skip this section.

- Step 1** Use a high-privilege account (with the USAGE permission for all sequences) to connect to the source database and run the following statement:

```
select 'SELECT pg_catalog.setval('||quote_literal(quote_ident(n.nspname))||','||quote_ident(c.relname))||','||nextval(c.oid)||');' as sqls from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname <> 'information_schema' and not (c.relname='hwdrs_ddl_info_id_seq' and n.nspname='public') order by n.nspname, c.relname;
```

The query result is the SQL statement that needs to be executed in the destination database.

- Step 2** Log in to the destination database as the destination database user used in the synchronization task and run the SQL statement queried in [Step 1](#) in the destination database.

- Step 3** Run the following statement in the destination database to check the sequence value synchronization result:

```
SELECT n.nspname, c.relname, nextval(c.oid) from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname <> 'information_schema' order by 1,2;
```

----End

Resetting Sequence Values in the Destination Database

If the source database failed and cannot be connected, you can reset the sequence values related to automatic increment or decrement columns in the destination database. If the source database can be connected, skip this section.

Step 1 Log in to the destination database as the destination database user used in the synchronization task.

Step 2 Run the following statement to query the SQL statement for resetting the sequence value corresponding to the sequence that uses nextval as the default value of the table column:

```
set search_path to "; select 'SELECT pg_catalog.setval('||quote_literal(quote_ident(s.sequence_schema))||'.||quote_ident(s.sequence_name)||', (SELECT '||case when s.increment::int<0 then 'min(' else 'max(' end||quote_ident(c.column_name)||')||case when s.increment::int<0 then '-1' else '+1' end||' FROM '||quote_ident(c.table_schema)||'.||quote_ident(c.table_name)||');' as sqls from information_schema.columns c join information_schema.sequences s on (position(quote_literal (quote_ident(s.sequence_schema))||'.||quote_ident(s.sequence_name)||':regclass' in c.column_default) > 0) where c.data_type in ('bigint', 'int', 'integer', 'smallint', 'numeric', 'real', 'double precision', 'double') and c.column_default like 'nextval(%%' order by s.sequence_schema, s.sequence_name;
```

The query result is the SQL statement that needs to be executed in the destination database.

Step 3 If the source database version is earlier than 10.0, skip this step. If the source database version is 10.0 or later, run the following statement in the destination database to query the SQL statement for resetting the sequence value corresponding to the additional column of the table identity column:

```
set search_path to "; select 'SELECT pg_catalog.setval('||quote_literal(seqname)||', (SELECT '||case when increment::int<0 then 'min(' else 'max(' end||colname||')||case when increment::int<0 then '-1' else '+1' end||' FROM '||tablename||');' as sqls from (select objid::regclass::text, refobjid::regclass::text, (pg_identify_object(refclassid,refobjid,refobjsubid)).identity, (pg_sequence_parameters(objid)).increment from pg_depend where deptype='i' and refobjsubid>0 and objid in (select c.oid from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind='S' and n.nspname !~ '^pg_' and n.nspname<>'information_schema')) p(seqname,tablename,colname,increment);
```

The query result is the SQL statement that needs to be executed in the destination database.

Step 4 Run the SQL statements queried in [Step 2](#) and [Step 3](#) in the destination database.

Step 5 Run the following statement in the destination database to check the sequence value synchronization result:

```
SELECT n.nspname, c.relname, nextval(c.oid) from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~ '^pg_' and n.nspname<>'information_schema' order by 1,2;
```

----End

13.6 Creating Event Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL

This section describes how to perform real-time synchronization from PostgreSQL to RDS for PostgreSQL. You can create event triggers and functions in the source database to obtain the DDL information of the source database, and then synchronize DDL operations to the destination database during the incremental synchronization phase.

Prerequisites

- The following DDL operations are supported:
 - Table-level synchronization: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), COMMENT ON COLUMN, and COMMENT ON TABLE
 - Database-level synchronization: TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA/TABLE, DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, ALTER VIEW, COMMENT ON COLUMN, COMMENT ON TABLE, COMMENT ON SCHEMA, COMMENT ON SEQUENCE, COMMENT ON INDEX, and COMMENT ON VIEW

CAUTION

- Table-level synchronization: If data is inserted into a renamed table, the data will not be synchronized to the destination database.
 - Database-level synchronization: Tables that are created not using the CREATE TABLE statement in the source database will not be synchronized to the destination database. For example, you run CREATE TABLE AS to create a table or call a function to create a table.
 - DDL statements starting with comments cannot be synchronized and are ignored.
 - DDL statements executed in functions and stored procedures cannot be synchronized and are ignored.
-
- If the source and destination databases are of different versions, use SQL statements that are compatible with both the source and destination databases to perform DDL operations. For example, if the source database is PostgreSQL 11 and the destination database is PostgreSQL 12, run the following statement to change the column type from char to int:

```
alter table tablename alter column columnname type int USING columnname::int;
```
 - Check whether a table named **hwdrs_ddl_info**, a function named **hwdrs_ddl_function()**, and a trigger named **hwdrs_ddl_event** exist in the source database in public mode. If they exist, delete them.
 - During database-level synchronization, if a table without a primary key is created, run the following command to set the replication attribute of the table without a primary key to full.

```
alter table tablename replica identity full;
```

Procedure

- If the source database is a self-managed PostgreSQL database or a database on another cloud platform, perform the following steps:
 - a. Connect to the database to be synchronized as a user who has permission to create event triggers.

- b. Run the following statements to create a table for storing DDL information:

```
DROP TABLE IF EXISTS public.hwdrs_ddl_info;
DROP SEQUENCE IF EXISTS public.hwdrs_ddl_info_id_seq;
CREATE TABLE public.hwdrs_ddl_info(
  id                bigserial primary key,
  ddl               text,
  username          varchar(64) default current_user,
  txid              varchar(16) default txid_current()::varchar(16),
  tag               varchar(64),
  database          varchar(64) default current_database(),
  schema            varchar(64) default current_schema,
  client_address    varchar(64) default inet_client_addr(),
  client_port       integer default inet_client_port(),
  event_time        timestamp default current_timestamp
);
```

- c. Run the following statements to create a function:

```
CREATE OR REPLACE FUNCTION public.hwdrs_ddl_function()
  RETURNS event_trigger
  LANGUAGE plpgsql
  SECURITY INVOKER
AS $$
  declare ddl text;
  declare real_num int;
  declare max_num int := 50000;
begin
  if (tg_tag in ('CREATE TABLE','ALTER TABLE','DROP TABLE','CREATE SCHEMA','CREATE SEQUENCE','ALTER SEQUENCE','DROP SEQUENCE','CREATE VIEW','ALTER VIEW','DROP VIEW','CREATE INDEX','ALTER INDEX','DROP INDEX','COMMENT')) then
    select current_query() into ddl;
    insert into public.hwdrs_ddl_info(ddl, username, txid, tag, database, schema, client_address, client_port, event_time)
      values (ddl, current_user, cast(txid_current() as varchar(16)), tg_tag, current_database(), current_schema, inet_client_addr(), inet_client_port(), current_timestamp);
    select count(id) into real_num from public.hwdrs_ddl_info;
    if real_num > max_num then
      if current_setting('server_version_num')::int < 100000 then
        delete from public.hwdrs_ddl_info where id < (select min(id)+1000 from public.hwdrs_ddl_info) and not exists (select 0 from pg_locks l join pg_database d on l.database=d.oid where d.datname=current_catalog and pid <> pg_backend_pid() and locktype='relation' and relation=to_regclass('public.hwdrs_ddl_info_pkey')::oid and mode='RowExclusiveLock');
      else
        delete from public.hwdrs_ddl_info where id < (select min(id)+1000 from public.hwdrs_ddl_info) and (xmax=0 or coalesce(txid_status(xmax::text::bigint), '') <> 'in progress');
      end if;
    end if;
  end if;
end;
$$;
```

- d. Run the following statements to grant necessary permissions to the objects created in **b** and **c**:

```
GRANT USAGE ON SCHEMA public TO public;
GRANT SELECT,INSERT,DELETE ON TABLE public.hwdrs_ddl_info TO public;
GRANT SELECT,USAGE ON SEQUENCE public.hwdrs_ddl_info_id_seq TO public;
GRANT EXECUTE ON FUNCTION public.hwdrs_ddl_function() TO public;
```

- e. Run the following statement to create a DDL event trigger:

```
CREATE EVENT TRIGGER hwdrs_ddl_event ON ddl_command_end EXECUTE PROCEDURE public.hwdrs_ddl_function();
```

- f. Run the following statement to set the created event trigger to enable:

```
ALTER EVENT TRIGGER hwdrs_ddl_event ENABLE ALWAYS;
```

- g. Return to the DRS console and create a PostgreSQL to RDS PostgreSQL synchronization task.

- h. After the synchronization task is complete, run the following statements to delete the created tables, functions, and triggers.
- ```
DROP EVENT trigger hwdrs_ddl_event;
DROP FUNCTION public.hwdrs_ddl_function();
DROP TABLE public.hwdrs_ddl_info;
```
- If the source database is RDS for PostgreSQL, perform the following steps:
    - a. Run the following statements to delete the created objects:

```
DROP EVENT TRIGGER IF EXISTS hwdrs_ddl_event;
DROP FUNCTION IF EXISTS public.hwdrs_ddl_function();
DROP TABLE IF EXISTS public.hwdrs_ddl_info;
```
    - b. Run the following statement as the **root** user to create a DDL plug-in:

```
SELECT CONTROL_EXTENSION('create', 'rds_hwdrs_ddl');
```
    - c. Run the following statements to update the function:

```
CREATE OR REPLACE FUNCTION public.hwdrs_ddl_function()
 RETURNS event_trigger
 LANGUAGE plpgsql
 SECURITY INVOKER
AS $BODY$
 declare ddl text;
 declare real_num int;
 declare max_num int := 50000;
begin
 if (tg_tag in ('CREATE TABLE','ALTER TABLE','DROP TABLE','CREATE SCHEMA','CREATE SEQUENCE','ALTER SEQUENCE','DROP SEQUENCE','CREATE VIEW','ALTER VIEW','DROP VIEW','CREATE INDEX','ALTER INDEX','DROP INDEX','COMMENT')) then
 select current_query() into ddl;
 insert into public.hwdrs_ddl_info(ddl, username, txid, tag, database, schema, client_address, client_port, event_time)
 values (ddl, current_user, cast(txid_current() as varchar(16)), tg_tag, current_database(), current_schema, inet_client_addr(), inet_client_port(), current_timestamp);
 select count(id) into real_num from public.hwdrs_ddl_info;
 if real_num > max_num then
 if current_setting('server_version_num')::int < 100000 then
 delete from public.hwdrs_ddl_info where id < (select min(id)+1000 from public.hwdrs_ddl_info) and not exists (select 0 from pg_locks l join pg_database d on l.database=d.oid where d.datname=current_catalog and pid <> pg_backend_pid() and locktype='relation' and relation=to_regclass('public.hwdrs_ddl_info_pkey')::oid and mode='RowExclusiveLock');
 else
 delete from public.hwdrs_ddl_info where id < (select min(id)+1000 from public.hwdrs_ddl_info) and (xmax=0 or coalesce(txid_status(xmax::text::bigint), '') <> 'in progress');
 end if;
 end if;
 end if;
end;
$BODY$;
```
    - d. Return to the DRS console and create a synchronization task from PostgreSQL to RDS for PostgreSQL.
    - e. After the synchronization task is complete, run the following statement to delete the created tables, functions, and triggers.

```
SELECT CONTROL_EXTENSION('drop', 'rds_hwdrs_ddl');
```



# 14 Appendix

---

## 14.1 General Usage Suggestions

Before using DRS, you need to understand the following:

- The destination database version must be the same as or later than the source database version.
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- To maintain data consistency before and after the synchronization, ensure that no data is written to your source and destination databases during a full synchronization. In the full+incremental synchronization mode, you can continue the synchronization while data is still being written to the source database.
- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.

- For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)
- You are advised to start your synchronization task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to be synchronized successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
- Data-Level Comparison  
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.
- If many-to-one synchronization is involved, you need to understand the following:
  - If you create a many-to-one synchronization task, refer to [Constraints and Operation Suggestions on Many-to-One Scenario](#).
  - For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail.
- After data synchronization is complete, you need to understand the following:
  - To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

## 14.2 DRS Pre-Check Items

A DRS synchronization process consists of four phases: task startup, full migration, incremental migration, and task completion. A single full migration task contains three phases. To ensure smooth synchronization, DRS automatically scans and checks parameters and objects in the source and destination databases before starting a synchronization task.

### Overview

View pre-check items based on the source database type of your synchronization task.

- The following lists the pre-check items for synchronization tasks with MySQL serving as the source.
  - [From MySQL to MySQL](#)
  - [From MySQL to PostgreSQL](#)
  - [MySQL->MariaDB](#)

## MySQL->MySQL

**Table 14-1** Pre-check items

| Category    | Check Item                       | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Solution to Failure                                                                                                                 |
|-------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Permissions | Source database permissions      | <ul style="list-style-type: none"> <li>Full synchronization requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT</li> <li>Full+incremental synchronization and incremental synchronization require the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</li> </ul>                                                                                                                                                       | For details, see <a href="#">Which MySQL Permissions Are Required for DRS?</a>                                                      |
|             | Destination database permissions | <ul style="list-style-type: none"> <li>The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES</li> </ul> <p>The <b>root</b> account of the RDS for MySQL DB instance has the preceding permissions by default.</p> <ul style="list-style-type: none"> <li>If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</li> </ul> | For details, see <a href="#">Which MySQL Permissions Are Required for DRS?</a>                                                      |
| Versions    | Source database versions         | Versions 5.5, 5.6, 5.7, and 8.0 are supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | For details, see <a href="#">Supported Databases.</a>                                                                               |
|             | Destination database versions    | Versions 5.5, 5.6, 5.7, and 8.0 are supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                              | For details, see <a href="#">Supported Databases.</a>                                                                               |
|             | Synchronization version          | The destination database version must be the same as or later than the source database version.                                                                                                                                                                                                                                                                                                                                                                                                             | For details, see <a href="#">Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version.</a> |

| Category   | Check Item                                       | Check Item Details                                                                                                                                                                                                                                                     | Solution to Failure                                                                                                            |
|------------|--------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Parameters | GTID status                                      | Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.                                                            | For details, see <a href="#">Checking Whether GTID Is Enabled for the Source Database</a> .                                    |
|            | Performance parameters                           | The <b>log_slave_updates</b> parameter of the source database must be enabled. Otherwise, the synchronization will fail.                                                                                                                                               | For details, see <a href="#">Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured</a> . |
|            |                                                  | The <b>binlog_row_image</b> parameter of the source database must be set to <b>FULL</b> . Otherwise, the synchronization will fail.                                                                                                                                    | For details, see <a href="#">Checking Whether the binlog_row_image Value is FULL</a> .                                         |
|            | Maximum size of a packet that can be transmitted | If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the <b>max_allowed_packet</b> value of the source database is too small, the synchronization task may fail.                                                          | For details, see <a href="#">Whether the max_allowed_packet Value of the Source Database Is Too Small</a> .                    |
|            |                                                  | If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the <b>max_allowed_packet</b> value of the destination database is too small, data cannot be written to the destination database and the full synchronization fails. | For details, see <a href="#">Checking Whether the max_allowed_packet Value of the Destination Database Is too Small</a> .      |
|            | sql_mode value                                   | If the MyISAM tables are included in the synchronization objects, the <b>sql_mode</b> parameter in the destination database cannot contain the <b>no_engine_substitution</b> parameter. Otherwise, the synchronization fails.                                          | For details, see <a href="#">Checking Whether the Source Database Contains Invalid sql_mode Values</a> .                       |

| Category                    | Check Item              | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Solution to Failure                                                                                           |
|-----------------------------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|
| Incremental synchronization | Index column length     | The index column length of the source database must meet requirements.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | For details, see <a href="#">Checking the Length of the Index Column in the Source Database</a> .             |
|                             | Binlog status           | During the incremental synchronization, the binlog of the source database must be enabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | For details, see <a href="#">Checking Whether the Source Database Binlog Is Enabled</a> .                     |
|                             | Binlog format           | The source database binlog must be row-based.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | For details, see <a href="#">Checking Whether the Source Database Binlog Is Row-Based</a> .                   |
|                             | Binlog retention period | <p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> <li>• If the source database is a self-managed MySQL database, set <b>expire_logs_days</b> to specify the binlog retention period. Set <b>expire_logs_day</b> 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.</li> <li>• If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in <a href="#">Setting a Local Retention Period for RDS for MySQL Binlogs</a>.</li> </ul> | For details, see <a href="#">Checking Whether the Binlog Retention Period Is Set on the Source Database</a> . |

| Category             | Check Item         | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                    | Solution to Failure                                                                                                                  |
|----------------------|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
|                      | server_id value    | <p>During an incremental synchronization, the <b>server_id</b> value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> <li>• If the source database version is MySQL 5.6 or earlier, the value of <b>server_id</b> ranges from 2 to 4294967296.</li> <li>• If the source database version is MySQL 5.7 or later, the value of <b>server_id</b> ranges from 1 to 4294967296.</li> </ul> | <p>For details, see <a href="#">Checking Whether the Source Database server_id Meets the Incremental Migration Requirements</a>.</p> |
|                      | Table fields       | <p>If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.</p>                                                                                                                                                                                                        | -                                                                                                                                    |
|                      | Additional columns | <p>In an incremental synchronization task, if no additional column is added to the destination database, the task fails.</p>                                                                                                                                                                                                                                                                                          | <p>For details, see <a href="#">Checking the Additional Column of the Destination Database</a>.</p>                                  |
| Destination database | Storage space      | <p>The destination DB instance must have sufficient storage space.</p>                                                                                                                                                                                                                                                                                                                                                | <p>For details, see <a href="#">Checking Whether the Destination Database Has Sufficient Storage Space</a>.</p>                      |
|                      | Status             | <p>The destination DB instance is running properly.</p>                                                                                                                                                                                                                                                                                                                                                               | -                                                                                                                                    |
| Consistency          | Character set      | <p>The character set of the destination database must be the same as that of the source database.</p>                                                                                                                                                                                                                                                                                                                 | <p>For details, see <a href="#">Checking Whether the Source and Destination Database Character Sets Are Consistent</a>.</p>          |

| Category | Check Item                                                                   | Check Item Details                                                                                                   | Solution to Failure                                                                                                                   |
|----------|------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
|          | Collation                                                                    | The <b>collation_server</b> value of the destination database must be the same as that of the source database.       | For details, see <a href="#">Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.</a>   |
|          | Clock                                                                        | The clock of the destination database must be the same as that of the source database.                               | -                                                                                                                                     |
|          | Time zone                                                                    | The <b>time_zone</b> value of the destination database must be the same as that of the source database.              | For details, see <a href="#">Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.</a>          |
|          | Case sensitive                                                               | The <b>lower_case_table_names</b> value of the destination database must be the same as that of the source database. | For details, see <a href="#">Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.</a> |
|          | Maximum length of the calculation result of the <b>group_concat</b> function | The <b>group_concat_max_len</b> value of the destination database must be the same as that of the source database.   | For details, see <a href="#">Checking Whether the Values of group_concat_max_len Are Consistent.</a>                                  |

| Category                | Check Item                      | Check Item Details                                                                                                                                                                                                                                                                                                                                                                | Solution to Failure                                                                                                                   |
|-------------------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
|                         | InnoDB check mode               | The <b>innodb_strict_mode</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                                                                                                                  | For details, see <a href="#">Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.</a> |
|                         | Data block encryption parameter | The <b>block_encryption_mode</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                                                                                                               | -                                                                                                                                     |
|                         | SQL mode                        | The <b>sql_mode</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                                                                                                                            | For details, see <a href="#">Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.</a>           |
| Synchronization objects | Selected objects                | <ul style="list-style-type: none"> <li>• Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized.</li> <li>• Only MyISAM and InnoDB tables can be synchronized.</li> <li>• Events and triggers cannot be synchronized.</li> <li>• Database parameters and the system database cannot be synchronized.</li> </ul> | -                                                                                                                                     |
|                         |                                 | 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.                                                                                                                                                       | For details, see <a href="#">Checking Database Mapping Objects.</a>                                                                   |



| Category | Check Item                  | Check Item Details                                                                                                                                                                                                                                                                                        | Solution to Failure                                                                                                                  |
|----------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
|          | Tables without primary keys | Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.                                                                                                                                                                      | For details, see <a href="#">Checking Whether the Source Database Tables Contain Primary Keys</a> .                                  |
|          | Associated objects          | Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys | For details, see <a href="#">Checking Whether Referenced Tables Are Selected for Migration</a> .                                     |
|          | Foreign key references      | Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.                                                                                             | For details, see <a href="#">Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database</a> . |
|          | Duplicate names             | If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.                                                                                            | For details, see <a href="#">Checking Whether the Names of the Source and Destination Databases Are the Same</a> .                   |

| Category           | Check Item                  | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Solution to Failure                                                                            |
|--------------------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
|                    | Database table naming rules | <ul style="list-style-type: none"> <li>• The source database names cannot contain non-ASCII characters, or the following characters: '&lt;&gt;/'</li> <li>• The source table and view names cannot contain non-ASCII characters, or the following characters: '&lt;&gt;/'</li> <li>• The source database name or mapped name cannot start with <b>ib_logfile</b> or be <b>ib_buffer_pool</b>, <b>ib_doublewrite</b>, <b>ibdata1</b> or <b>ibtmp1</b>.</li> <li>• The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (&lt;&gt;), backslash (\), and single quotation marks (')</li> <li>• 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.</li> </ul> | -                                                                                              |
| SSL                | SSL configuration           | If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | For details, see <a href="#">Checking Whether the SSL Connection Is Correctly Configured</a> . |
| Network conditions | Network conditions          | The IP address, port, username, and password of the destination database are correctly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | For details, see <a href="#">Checking Whether the Destination Database Is Connected</a> .      |
|                    |                             | The IP address, port, username, and password of the source database are correctly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | For details, see <a href="#">Checking Whether the Source Database Is Connected</a> .           |

## MySQL->PostgreSQL

Table 14-2 Pre-check items

| Category    | Check Item                                       | Check Item Details                                                                                                                                                                                                                                                            | Solution to Failure                                                                                         |
|-------------|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| Permissions | Source database permissions                      | <ul style="list-style-type: none"> <li>Full synchronization requires the following minimum permission: SELECT</li> <li>Full+incremental synchronization requires the following minimum permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</li> </ul> | For details, see <a href="#">Which MySQL Permissions Are Required for DRS?</a>                              |
|             | Destination database permissions                 | The account of the RDS for PostgreSQL instance has the permissions by default.                                                                                                                                                                                                | For details, see <a href="#">Which MySQL Permissions Are Required for DRS?</a>                              |
| Versions    | Source database versions                         | Versions 5.5, 5.6, 5.7, and 8.0 are supported.                                                                                                                                                                                                                                | For details, see <a href="#">Supported Databases</a> .                                                      |
|             | Destination database versions                    | Versions 9.5, 9.6, 10, and 11 are supported.                                                                                                                                                                                                                                  | For details, see <a href="#">Supported Databases</a> .                                                      |
| Parameters  | GTID status                                      | Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.                                                                   | For details, see <a href="#">Checking Whether GTID Is Enabled for the Source Database</a> .                 |
|             | Performance parameters                           | The <b>binlog_row_image</b> parameter of the source database must be set to <b>FULL</b> . Otherwise, the synchronization will fail.                                                                                                                                           | For details, see <a href="#">Checking Whether the binlog_row_image Value is FULL</a> .                      |
|             | Maximum size of a packet that can be transmitted | If there is a lot of data to be synchronized or there are too many fields to be synchronized, and the <b>max_allowed_packet</b> value of the source database is too small, the synchronization task may fail.                                                                 | For details, see <a href="#">Whether the max_allowed_packet Value of the Source Database Is Too Small</a> . |

| Category                    | Check Item      | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                                                               | Solution to Failure                                                                                                           |
|-----------------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Incremental synchronization | Binlog status   | During the incremental synchronization, the binlog of the source database must be enabled.                                                                                                                                                                                                                                                                                                                                                                       | For details, see <a href="#">Checking Whether the Source Database Binlog Is Enabled.</a>                                      |
|                             | Binlog format   | The source database binlog must be row-based.                                                                                                                                                                                                                                                                                                                                                                                                                    | For details, see <a href="#">Checking Whether the Source Database Binlog Is Row-Based.</a>                                    |
|                             | server_id value | During an incremental synchronization, the <b>server_id</b> value of the MySQL source database must be set. <ul style="list-style-type: none"> <li>If the source database version is MySQL 5.6 or earlier, the value of <b>server_id</b> ranges from 2 to 4294967296.</li> <li>If the source database version is MySQL 5.7 or later, the value of <b>server_id</b> ranges from 1 to 4294967296.</li> </ul>                                                       | For details, see <a href="#">Checking Whether the Source Database server_id Meets the Incremental Migration Requirements.</a> |
|                             | Table fields    | <ul style="list-style-type: none"> <li>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.</li> <li>The following table field types are not supported in the source database: xml, geometry, point, lineString, polygon, geometrycollection, multipoint, multilinestring, multipolygon, and json.</li> </ul> | -                                                                                                                             |
| Destination database        | Storage space   | The destination DB instance must have sufficient storage space.                                                                                                                                                                                                                                                                                                                                                                                                  | For details, see <a href="#">Checking Whether the Destination Database Has Sufficient Storage Space.</a>                      |

| Category                | Check Item                  | Check Item Details                                                                                                                                                                                                                                                        | Solution to Failure                                                                                                          |
|-------------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
|                         | Status                      | The destination DB instance is running properly.                                                                                                                                                                                                                          | -                                                                                                                            |
| Consistency             | Character set               | The character set of the destination database must be the same as that of the source database.                                                                                                                                                                            | For details, see <a href="#">Checking Whether the Source and Destination Database Character Sets Are Consistent.</a>         |
|                         | Time zone                   | The <b>time_zone</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                   | For details, see <a href="#">Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.</a> |
|                         | Table structures            | The table structure of the destination database must be the same as that of the source database.                                                                                                                                                                          | For details, see <a href="#">Checking Whether Tables Structures Are Consistent.</a>                                          |
| Synchronization objects | Selected objects            | <ul style="list-style-type: none"> <li>• Only table structures, table data, and indexes can be synchronized.</li> <li>• Only MyISAM and InnoDB tables can be synchronized.</li> <li>• Other database objects such as stored procedures cannot be synchronized.</li> </ul> | -                                                                                                                            |
|                         | Tables without primary keys | Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.                                                                                                                                      | For details, see <a href="#">Checking Whether the Source Database Tables Contain Primary Keys.</a>                           |
|                         | Primary keys                | Do not use an imprecise value type as the primary key in the database. This feature affects the synchronization of UPDATE and DELETE statements in the DRS incremental scenario.                                                                                          | -                                                                                                                            |

| Category           | Check Item                       | Check Item Details                                                                                                                                                                                            | Solution to Failure                                                                                                                 |
|--------------------|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
|                    | Foreign key reference operations | Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. | For details, see <a href="#">Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.</a> |
|                    | Duplicate names                  | The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).                                                                           | For details, see <a href="#">Checking Whether the Names of the Source and Destination Databases Are the Same.</a>                   |
|                    | Database table naming rules      | The names of the source databases and tables cannot contain non-ASCII characters, or the following characters: '<>/'                                                                                          | -                                                                                                                                   |
|                    | Object name length               | The destination database object name can contain a maximum of 63 characters.                                                                                                                                  | For details, see <a href="#">Checking Whether the Length of the Source Database Object Names Exceeds the Limit.</a>                 |
|                    | Source database                  | The source database is properly connected during the synchronization object check in the pre-check phase.                                                                                                     | For details, see <a href="#">Checking the Synchronization Objects.</a>                                                              |
| SSL                | SSL configuration                | If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.                                                                                                | For details, see <a href="#">Checking Whether the SSL Connection Is Correctly Configured.</a>                                       |
| Network conditions | Network conditions               | The IP address, port, username, and password of the destination database are correctly configured.                                                                                                            | For details, see <a href="#">Checking Whether the Destination Database Is Connected.</a>                                            |

| Category | Check Item | Check Item Details                                                                            | Solution to Failure                                                                 |
|----------|------------|-----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
|          |            | The IP address, port, username, and password of the source database are correctly configured. | For details, see <a href="#">Checking Whether the Source Database Is Connected.</a> |

## MySQL->MariaDB

**Table 14-3** Pre-check items

| Category    | Check Item                       | Check Item Details                                                                                                                                                                                                                                                                                                                                    | Solution to Failure                                                            |
|-------------|----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| Permissions | Source database permissions      | <ul style="list-style-type: none"> <li>Full synchronization requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT</li> <li>Full+incremental synchronization and incremental synchronization require the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</li> </ul> | For details, see <a href="#">Which MySQL Permissions Are Required for DRS?</a> |
|             | Destination database permissions | <ul style="list-style-type: none"> <li>The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, REFERENCES, and INDEX</li> </ul> <p>The <b>root</b> account of the RDS for MariaDB instance has the preceding permissions by default.</p>                 | For details, see <a href="#">Which MySQL Permissions Are Required for DRS?</a> |
| Versions    | Source database versions         | Versions 5.5, 5.6, 5.7, and 8.0 are supported.                                                                                                                                                                                                                                                                                                        | For details, see <a href="#">Supported Databases.</a>                          |
|             | Destination database versions    | Version 10.5 is supported.                                                                                                                                                                                                                                                                                                                            | For details, see <a href="#">Supported Databases.</a>                          |

| Category                    | Check Item             | Check Item Details                                                                                                                                                                                                            | Solution to Failure                                                                                                            |
|-----------------------------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Parameters                  | GTID status            | Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.                   | For details, see <a href="#">Checking Whether GTID Is Enabled for the Source Database</a> .                                    |
|                             | Performance parameters | The <b>log_slave_updates</b> parameter of the source database must be enabled. Otherwise, the synchronization will fail.                                                                                                      | For details, see <a href="#">Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured</a> . |
|                             |                        | The <b>binlog_row_image</b> parameter of the source database must be set to <b>FULL</b> . Otherwise, the synchronization will fail.                                                                                           | For details, see <a href="#">Checking Whether the binlog_row_image Value is FULL</a> .                                         |
|                             | sql_mode value         | If the MyISAM tables are included in the synchronization objects, the <b>sql_mode</b> parameter in the destination database cannot contain the <b>no_engine_substitution</b> parameter. Otherwise, the synchronization fails. | For details, see <a href="#">Checking Whether the Source Database Contains Invalid sql_mode Values</a> .                       |
| Incremental synchronization | Index column length    | The index column length of the source database must meet requirements.                                                                                                                                                        | For details, see <a href="#">Checking the Length of the Index Column in the Source Database</a> .                              |
|                             | Binlog status          | During the incremental synchronization, the binlog of the source database must be enabled.                                                                                                                                    | For details, see <a href="#">Checking Whether the Source Database Binlog Is Enabled</a> .                                      |
|                             | Binlog format          | The source database binlog must be row-based.                                                                                                                                                                                 | For details, see <a href="#">Checking Whether the Source Database Binlog Is Row-Based</a> .                                    |



| Category | Check Item              | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | Solution to Failure                                                                                                            |
|----------|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
|          | Binlog retention period | <p>If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the synchronization may fail.</p> <ul style="list-style-type: none"> <li>• If the source database is a self-managed MySQL database, set <b>expire_logs_days</b> to specify the binlog retention period. Set <b>expire_logs_day</b> 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.</li> <li>• If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in <a href="#">Setting a Local Retention Period for RDS for MySQL Binlogs</a>.</li> </ul> | For details, see <a href="#">Checking Whether the Binlog Retention Period Is Set on the Source Database</a> .                  |
|          | server_id value         | <p>During an incremental synchronization, the <b>server_id</b> value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> <li>• If the source database version is MySQL 5.6 or earlier, the value of <b>server_id</b> ranges from 2 to 4294967296.</li> <li>• If the source database version is MySQL 5.7 or later, the value of <b>server_id</b> ranges from 1 to 4294967296.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                            | For details, see <a href="#">Checking Whether the Source Database server_id Meets the Incremental Migration Requirements</a> . |
|          | Table fields            | <p>If there is a table containing fields of the longtext or longblob type in the synchronization object, you are advised to create a DRS task with large specifications. Otherwise, capture OOM may occur.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | -                                                                                                                              |
|          | Additional columns      | <p>In an incremental synchronization task, if no additional column is added to the destination database, the task fails.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | For details, see <a href="#">Checking the Additional Column of the Destination Database</a> .                                  |

| Category             | Check Item     | Check Item Details                                                                                                   | Solution to Failure                                                                                                                   |
|----------------------|----------------|----------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Destination database | Storage space  | The destination DB instance must have sufficient storage space.                                                      | For details, see <a href="#">Checking Whether the Destination Database Has Sufficient Storage Space.</a>                              |
|                      | Status         | The destination DB instance is running properly.                                                                     | -                                                                                                                                     |
| Consistency          | Character set  | The character set of the destination database must be the same as that of the source database.                       | For details, see <a href="#">Checking Whether the Source and Destination Database Character Sets Are Consistent.</a>                  |
|                      | Collation      | The <b>collation_server</b> value of the destination database must be the same as that of the source database.       | For details, see <a href="#">Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.</a>   |
|                      | Clock          | The clock of the destination database must be the same as that of the source database.                               | -                                                                                                                                     |
|                      | Case sensitive | The <b>lower_case_table_names</b> value of the destination database must be the same as that of the source database. | For details, see <a href="#">Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.</a> |

| Category                | Check Item                                                                   | Check Item Details                                                                                                                                                                                                                                                                                                                                                                | Solution to Failure                                                                                                                   |
|-------------------------|------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
|                         | Maximum length of the calculation result of the <b>group_concat</b> function | The <b>group_concat_max_len</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                                                                                                                | For details, see <a href="#">Checking Whether the Values of group_concat_max_len Are Consistent.</a>                                  |
|                         | InnoDB check mode                                                            | The <b>innodb_strict_mode</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                                                                                                                  | For details, see <a href="#">Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.</a> |
|                         | SQL mode                                                                     | The <b>sql_mode</b> value of the destination database must be the same as that of the source database.                                                                                                                                                                                                                                                                            | For details, see <a href="#">Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.</a>           |
| Synchronization objects | Selected objects                                                             | <ul style="list-style-type: none"> <li>• Tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized.</li> <li>• Only MyISAM and InnoDB tables can be synchronized.</li> <li>• Events and triggers cannot be synchronized.</li> <li>• Database parameters and the system database cannot be synchronized.</li> </ul> | -                                                                                                                                     |
|                         |                                                                              | 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.                                                                                                                                                       | For details, see <a href="#">Checking Database Mapping Objects.</a>                                                                   |

| Category | Check Item                  | Check Item Details                                                                                                                                                                                                                                                                                        | Solution to Failure                                                                                                                 |
|----------|-----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
|          | Tables without primary keys | Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.                                                                                                                                                                      | For details, see <a href="#">Checking Whether the Source Database Tables Contain Primary Keys.</a>                                  |
|          | Associated objects          | Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys | For details, see <a href="#">Checking Whether Referenced Tables Are Selected for Migration.</a>                                     |
|          | Foreign key references      | Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.                                                                                             | For details, see <a href="#">Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.</a> |
|          | Duplicate names             | If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.                                                                                            | For details, see <a href="#">Checking Whether the Names of the Source and Destination Databases Are the Same.</a>                   |

| Category           | Check Item                  | Check Item Details                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Solution to Failure                                                                            |
|--------------------|-----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
|                    | Database table naming rules | <ul style="list-style-type: none"> <li>• The source database names cannot contain non-ASCII characters, or the following characters: '&lt;&gt;\'\"</li> <li>• The source table and view names cannot contain non-ASCII characters, or the following characters: '&lt;&gt;\'\"</li> <li>• The source database name or mapped name cannot start with <b>ib_logfile</b> or be <b>ib_buffer_pool</b>, <b>ib_doublewrite</b>, <b>ibdata1</b> or <b>ibtmp1</b>.</li> <li>• The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (&lt;&gt;), backslash (\), and single quotation marks (')</li> <li>• 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.</li> </ul> | -                                                                                              |
| SSL                | SSL configuration           | If you enable SSL, the SSL connections for both the source and destination databases are correctly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | For details, see <a href="#">Checking Whether the SSL Connection Is Correctly Configured</a> . |
| Network conditions | Network conditions          | The IP address, port, username, and password of the destination database are correctly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | For details, see <a href="#">Checking Whether the Destination Database Is Connected</a> .      |
|                    |                             | The IP address, port, username, and password of the source database are correctly configured.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | For details, see <a href="#">Checking Whether the Source Database Is Connected</a> .           |

# A Change History

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2024-02-28  | <p>This issue is the fifty-seventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Supported full synchronization and incremental synchronization for real-time synchronization from Microsoft SQL Server to GaussDB.</li><li>• Supported DDL type selection for real-time synchronization from GaussDB Primary/Standby to Oracle.</li><li>• Supported a DR cluster serving as the source database for a self-built synchronization task from GaussDB Primary/Standby.</li><li>• Supported sharding mode selection for real-time synchronization from Oracle.</li><li>• Supported LOB length comparison in value comparison for real-time synchronization from Oracle to GaussDB.</li></ul> |
| 2023-12-30  | <p>This issue is the fifty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Supported task parameter changing for real-time synchronization.</li><li>• Supported real-time synchronization from MySQL to MariaDB.</li><li>• Supported real-time synchronization from MariaDB to MySQL.</li><li>• Supported real-time synchronization from MariaDB to GaussDB(for MySQL).</li><li>• Supported object file importing, flow control, task pausing and task resetting for real-time synchronization from PostgreSQL to GaussDB.</li></ul>                                                                                                                                                  |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2023-11-30  | <p>This issue is the fifty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>● Supported ultra-large specifications for real-time synchronization.</li> <li>● Supported progress details viewing for real-time synchronization.</li> <li>● Supported PostgreSQL 15 for DRS real-time synchronization.</li> <li>● Supported real-time synchronization from Microsoft SQL Server to Kafka.</li> <li>● Supported real-time synchronization from OceanBase (MySQL-compatible) to GaussDB(for MySQL).</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>● The following data flow scenarios meet the commercial use standards. <ul style="list-style-type: none"> <li>- Real-time synchronization from MySQL to GaussDB Primary/Standby</li> <li>- Real-time synchronization from MySQL to GaussDB Distributed</li> <li>- Real-time synchronization from DDS to Kafka</li> <li>- Real-time synchronization from GaussDB Primary/Standby to MySQL</li> <li>- Real-time synchronization from GaussDB Primary/Standby to GaussDB Distributed</li> <li>- Real-time synchronization from GaussDB Distributed to MySQL</li> <li>- Real-time synchronization from GaussDB Distributed to GaussDB Primary/Standby</li> </ul> </li> </ul> |
| 2023-10-30  | <p>This issue is the fifty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>● Supported partition table synchronization in the incremental phase for real-time synchronization with DB2 for LUW serving as the source database.</li> <li>● Added the requirement for uploading the JDBC driver package when testing the connection for real-time synchronization with DB2 for LUW serving as the source database.</li> <li>● Added support for upgrading task specifications in a DRS multi-specification task.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 2023-09-30  | <p>This issue is the fifty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>● Supported synchronization position resetting for tasks from MySQL to GaussDB(DWS).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2023-08-30  | <p>This issue is the fifty-second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported customizing DROP PARTITION and RENAME COLUMN for incremental DDL synchronization from MySQL to MySQL.</li> <li>• Supported COMMENT ON for incremental DDL synchronization from PostgreSQL to PostgreSQL.</li> <li>• Supported DRS task filtering by DB instance ID or database IP address.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                  |
| 2023-07-30  | <p>This issue is the fifty-first official release, which incorporates the following change:</p> <ul style="list-style-type: none"> <li>• Supported AZ selection for DRS synchronization tasks.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| 2023-06-30  | <p>This issue is the fiftieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported full synchronization and incremental synchronization for real-time synchronization from DDM to MySQL and DDM to Oracle.</li> <li>• Supported the selection of synchronization object types for real-time synchronization from MySQL to GaussDB.</li> <li>• Supported the object name mapping for real-time synchronization from GaussDB to GaussDB.</li> <li>• Supported database-level synchronization and schema-level synchronization for real-time synchronization from GaussDB to Kafka.</li> <li>• Supported data filtering for real-time synchronization from DB2 for LUW to GaussDB.</li> <li>• Supported Oracle 21c for DRS real-time synchronization.</li> </ul> |
| 2023-05-30  | <p>This issue is the forty-ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported task import for creating synchronization tasks in batches.</li> <li>• Supported out-of-cloud synchronization from DDS to Kafka.</li> <li>• Supported GaussDB(DWS) 8.2.0.</li> <li>• Supported DDL type selection for synchronization tasks from Oracle to GaussDB.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                           |



| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2023-04-30  | <p>This issue is the forty-eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported real-time synchronization of DDS 4.2.</li> <li>• Supported specification upgrade in single-node MySQL-to-Kafka synchronization tasks.</li> <li>• Supported the selection of synchronization object types for synchronization from MySQL to Kafka.</li> <li>• Supported full data read from the standby database for synchronization from GaussDB(for MySQL) to GaussDB(for MySQL).</li> <li>• Supported the full mode for to-the-cloud synchronization from MySQL to MySQL.</li> <li>• Supported the incremental mode for real-time synchronization from Oracle to MySQL and from Oracle to GaussDB(DWS).</li> <li>• Supported the full and full+incremental modes for real-time synchronization from GaussDB Distributed to MySQL.</li> <li>• Supported PostgreSQL 13 and 14 for to-the-cloud real-time synchronization from Oracle to PostgreSQL.</li> <li>• Supported quick diagnosis if a DRS connection test fails.</li> <li>• Supported MariaDB synchronization (including to and from the cloud).</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• The following data flow scenarios meet the commercial use standards. <ul style="list-style-type: none"> <li>– Real-time synchronization from Oracle to GaussDB Primary/Standby</li> <li>– Real-time synchronization from Oracle to GaussDB Distributed</li> <li>– Real-time synchronization from GaussDB(for MySQL) to GaussDB(for MySQL)</li> <li>– Real-time synchronization from GaussDB Primary/Standby to GaussDB Primary/Standby</li> <li>– Real-time synchronization from GaussDB Distributed to GaussDB Distributed</li> <li>– Real-time synchronization from GaussDB(for MySQL) to GaussDB(for MySQL)</li> </ul> </li> <li>• The following data flow scenarios are in the open beta test phase. <ul style="list-style-type: none"> <li>– Real-time synchronization from MySQL to GaussDB Primary/Standby</li> <li>– Real-time synchronization from DB2 for LUW to GaussDB Primary/Standby</li> </ul> </li> </ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|             | <ul style="list-style-type: none"> <li>- Real-time synchronization from DB2 for LUW to GaussDB Distributed</li> <li>- Real-time synchronization from Microsoft SQL Server to GaussDB Primary/Standby</li> <li>- Real-time synchronization from Microsoft SQL Server to GaussDB Distributed</li> <li>- Real-time synchronization from Microsoft SQL Server to GaussDB(DWS)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                     |
| 2023-03-30  | <p>This issue is the forty-seventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported to-the-cloud synchronization from GaussDB(for MySQL) to GaussDB(for MySQL).</li> <li>• Supported out-of-cloud synchronization from PostgreSQL to PostgreSQL.</li> <li>• Supported task type selection for synchronization from GaussDB(for MySQL) to CSS/ES.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• On the DRS task creation page, changed <b>Single</b> and <b>Primary/Standby</b> to <b>Single-AZ</b> and <b>Dual-AZ</b> in the <b>DRS Task Type</b> area.</li> </ul>                                                                         |
| 2023-02-28  | <p>This issue is the forty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported the sorting of row comparison results in ascending or descending order by <b>Source Database Table Rows</b> or <b>Destination Database Table Rows</b>.</li> <li>• Supported specification upgrade in single-node MySQL-to-GaussDB(DWS) synchronization tasks.</li> </ul>                                                                                                                                                                                                                                                                                                                             |
| 2023-01-30  | <p>This issue is the forty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported PostgreSQL 14 for real-time synchronization with PostgreSQL serving as the source.</li> <li>• Supported DB2SECURITYLABEL data for real-time synchronization with DB2 for LUW serving as the source.</li> <li>• Supported Elasticsearch 7.10 for real-time synchronization from MySQL to CSS/ES and from GaussDB(for MySQL) to CSS/ES.</li> <li>• Supported start point setting for incremental synchronization from MySQL to Kafka and from GaussDB(for MySQL) to Kafka.</li> <li>• Supported real-time synchronization from Microsoft SQL Server to Microsoft SQL Server (to the cloud).</li> </ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2022-12-30  | <p>This issue is the forty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Supported multiple specifications for real-time synchronization from MySQL to CSS/ES.</li><li>• Supported the topic and partition policies for real-time synchronization from Oracle to Kafka.</li><li>• A synchronization task can be created by specifying DN connection information when the source is self-built GaussDB distributed.</li><li>• Supported the full+incremental mode for real-time synchronization from MySQL to Kafka and synchronization from GaussDB(for MySQL) to Kafka.</li><li>• Supported monitoring information viewing for real-time synchronization from MySQL to Kafka and synchronization from GaussDB to Kafka.</li></ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2022-11-30  | <p>This issue is the forty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• After DRS interconnects with LTS and log reporting to LTS is enabled, all logs generated by DRS instances will be uploaded to LTS for management.</li> <li>• Supported real-time synchronization from Microsoft SQL Server to GaussDB primary/standby (to the cloud)</li> <li>• Supported real-time synchronization from Microsoft SQL Server to GaussDB distributed (to the cloud)</li> <li>• Supported specification upgrade in single-node MySQL synchronization tasks.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• The following data flow scenarios meet the commercial use standards. <ul style="list-style-type: none"> <li>– Real-time synchronization from Oracle to GaussDB primary/standby</li> <li>– Real-time synchronization from Oracle to GaussDB distributed</li> <li>– Real-time synchronization from MongoDB to DDS</li> <li>– Real-time synchronization from GaussDB primary/standby to Oracle</li> <li>– Real-time synchronization from GaussDB primary/standby to GaussDB(DWS)</li> <li>– Real-time synchronization from GaussDB primary/standby to Kafka</li> <li>– Real-time synchronization from GaussDB distributed to Oracle</li> <li>– Real-time synchronization from GaussDB distributed to GaussDB(DWS)</li> <li>– Real-time synchronization from GaussDB distributed to Kafka</li> </ul> </li> </ul> |
| 2022-10-30  | <p>This issue is the forty-second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Added support for selecting the synchronization object type during real-time synchronization from Microsoft SQL Server to GaussDB (DWS).</li> <li>• Added the GTID status pre-check for the source database when MySQL is used as the source in synchronization.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2022-08-30  | <p>This issue is the forty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• The WE8ISO8859P15 character set is supported when the source database is an Oracle database for data synchronization.</li> <li>• Real-time synchronization from PostgreSQL to Kafka supports multiple specifications.</li> <li>• Abnormal data policies can be set for data synchronization from MySQL to GaussDB(DWS).</li> </ul>                                                                                                                                                                                                                              |
| 2022-07-30  | <p>This issue is the fortieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Added verification of out-of-line storage of primary key columns when PostgreSQL is the source database for data synchronization.</li> <li>• Added support for synchronizing different table objects to different topics when <b>Synchronization Object</b> is set to <b>Import object file</b> for Oracle to Kafka synchronization.</li> <li>• Supported synchronization of MySQL JSON data to GaussDB(DWS).</li> <li>• If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically stops.</li> </ul> |
| 2022-06-30  | <p>This issue is the thirty-ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported periodic row comparison for real-time synchronization from MySQL to GaussDB(DWS).</li> <li>• Added support for skipping DDL operations during real-time synchronization from MySQL to GaussDB(DWS).</li> </ul>                                                                                                                                                                                                                                                                                                                                       |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2022-05-30  | <p>This issue is the thirty-eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>● Real-time synchronization from MongoDB to DDS (to the cloud)</li> <li>● Real-time synchronization from Microsoft SQL Server to GaussDB(DWS) (to the cloud)</li> <li>● Real-time synchronization from TiDB to GaussDB(for MySQL)</li> <li>● Real-time synchronization from DB2 for LUW to GaussDB(DWS) (to the cloud)</li> <li>● Supported multiple specifications for real-time synchronization from GaussDB primary/standby to Kafka and from GaussDB ditributed to Kafka.</li> <li>● Supported column mapping for real-time synchronization from MySQL to CSS/ES.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>● The following meets the commercial use standards. <ul style="list-style-type: none"> <li>– Real-time synchronization from DDM to Oracle</li> <li>– Real-time synchronization from MySQL to Oracle</li> <li>– Real-time synchronization from GaussDB(for MySQL) to MySQL</li> <li>– Real-time synchronization from GaussDB(for MySQL) to GaussDB(DWS)</li> <li>– Real-time synchronization from GaussDB(for MySQL) to Kafka</li> <li>– Real-time synchronization from GaussDB(for MySQL) to Oracle</li> <li>– Real-time synchronization from GaussDB(for MySQL) to CSS/ES</li> <li>– Real-time synchronization from PostgreSQL to Kafka</li> </ul> </li> </ul> |
| 2022-04-30  | <p>This issue is the thirty-seventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>● Supported task cloning for real-time synchronization from MySQL to CSS/ES.</li> <li>● Supported task cloning for real-time synchronization from PostgreSQL to PostgreSQL.</li> <li>● Supported SSL for the source database that is DB2 for LUW 10.5 or later.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>● Adjusted the length and character range of tag keys and tag values.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2022-03-30  | <p>This issue is the thirty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported data filtering and column filtering for real-time synchronization from MySQL to CSS/ES and from GaussDB(for MySQL) to CSS/ES.</li> <li>• Supported single incremental mode for real-time synchronization from MySQL to GaussDB(DWS).</li> <li>• Supported DB2 for LUW 10.1 and 11.1 as the source and GaussDB as the destination during synchronization.</li> <li>• Supported Oracle as the source in a real-time synchronization task. The accumulated number of operations on each table can be displayed.</li> <li>• Supported suspending, resetting, and cloning real-time synchronization tasks from DDS to MongoDB.</li> <li>• Supported real-time out-of-cloud synchronization from GaussDB distributed to GaussDB primary/standby.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• Changed the length of the WHERE condition for data filtering to 512 characters.</li> <li>• Supported disable task delay notification.</li> </ul> |
| 2022-02-28  | <p>This issue is the thirty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported specifying the replication slot name for real-time synchronization from GaussDB to Kafka and from PostgreSQL to Kafka.</li> <li>• Supported exporting snapshots during the synchronization of PostgreSQL databases.</li> <li>• Supported multiple specifications for some real-time synchronization tasks.</li> <li>• Supported stopping tasks in batches.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 2022-01-30  | <p>This issue is the thirty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Real-time synchronization from DB2 for LUW to GaussDB supported DB2 for LUW 10.5 and 11.5 and large objects.</li> <li>• Supported PostgreSQL to Kafka synchronization (out of the cloud and self-built to self-built)</li> <li>• Supported importing objects from files during PostgreSQL to PostgreSQL synchronization.</li> <li>• Supported SSL for synchronization from MySQL to CSS/ES and GaussDB(for MySQL) to CSS/ES.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• Put MySQL to CSS/ES synchronization into commercial user.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                    |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2021-12-31  | <p>This issue is the thirty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added the description about the impact of DRS on databases.</li><li>• Supported DDLs during PostgreSQL incremental synchronization.</li><li>• Supported JSON during synchronization from GaussDB to Kafka.</li><li>• Supported rate limiting for real-time synchronization from MySQL to GaussDB(for MySQL).</li><li>• Supported real-time synchronization from MySQL to GaussDB primary/standby in the cloud.</li><li>• DRS supports real-time synchronization from GaussDB primary/standby to MySQL.</li><li>• Supported real-time synchronization tasks billed on a yearly/monthly basis.</li><li>• Supported Oracle-to-DDM synchronization in the scenario where the destination has more columns than the source.</li><li>• Supported pausing and resetting MySQL to GaussDB(DWS) synchronization tasks.</li><li>• Supported pausing and resetting PostgreSQL to PostgreSQL synchronization tasks.</li></ul> <p>Changed the following content:</p> <ul style="list-style-type: none"><li>• Moved the Send Notifications option to the task confirmation page.</li><li>• The real-time synchronization of GaussDB as the source did not support database-level synchronization.</li></ul> |



| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2021-11-30  | <p>This issue is the thirty-second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported real-time synchronization from GaussDB primary/standby to GaussDB(DWS).</li> <li>• Supported real-time synchronization from GaussDB(for MySQL) to MySQL.</li> <li>• Supported real-time synchronization from GaussDB(for MySQL) to GaussDB(DWS).</li> <li>• Supported real-time synchronization from DB2 for LUW to GaussDB primary/standby.</li> <li>• Supported real-time synchronization from DB2 for LUW to GaussDB distributed.</li> <li>• Supported re-editing and suspending real-time synchronization from PostgreSQL to GaussDB(DWS).</li> <li>• Supported real-time synchronization from GaussDB distributed to GaussDB(DWS) in full+incremental mode.</li> <li>• Supported adding field types to additional columns during MySQL to GaussDB(for MySQL) data processing.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• The following scenarios are in the open beta test phase. <ul style="list-style-type: none"> <li>- Real-time synchronization from PostgreSQL to GaussDB(DWS)</li> <li>- Real-time synchronization from DDM to MongoDB</li> <li>- Real-time synchronization from GaussDB primary/standby to Oracle</li> <li>- Real-time synchronization from GaussDB primary/standby to Kafka</li> <li>- Supported real-time synchronization from GaussDB primary/standby to GaussDB distributed.</li> <li>- Supported real-time synchronization from GaussDB primary/standby to GaussDB primary/standby.</li> <li>- Real-time synchronization from GaussDB distributed to MySQL</li> <li>- Real-time synchronization from GaussDB(for MySQL) to Kafka</li> <li>- Real-time synchronization from GaussDB(for MySQL) to Oracle</li> <li>- Real-time synchronization from GaussDB(for MySQL) to CSS/ES</li> <li>- Real-time synchronization from self-built MySQL to CSS/ES</li> </ul> </li> </ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2021-09-30  | <p>This issue is the thirty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Added the description of the product architecture and principles.</li> <li>• Supported real-time synchronization from DDM to DDM.</li> <li>• DRS supports real-time synchronization from MySQL to Oracle.</li> <li>• Supported real-time synchronization from GaussDB(for MySQL) to Oracle.</li> <li>• Supported real-time synchronization from GaussDB(for MySQL) to Elasticsearch.</li> <li>• DRS supports real-time incremental synchronization of MongoDB.</li> <li>• Supported real-time synchronization from self-built MySQL to Elasticsearch.</li> <li>• Supported object-level and row-level comparisons during MySQL to GaussDB distributed real-time synchronization.</li> <li>• Supported DDL filtering in the following scenarios: MySQL to MySQL, MySQL to GaussDB(for MySQL), MySQL to GaussDB(DWS), and MySQL to PostgreSQL synchronization.</li> <li>• Supported task cloning for MySQL real-time synchronization.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• The following scenarios meet the commercial user standard. <ul style="list-style-type: none"> <li>- Real-time synchronization from Oracle to DDM</li> <li>- Real-time synchronization from Oracle to PostgreSQL</li> <li>- Real-time synchronization from DDM to Kafka</li> <li>- Real-time synchronization from DDM to GaussDB(DWS)</li> <li>- Real-time synchronization from PostgreSQL to PostgreSQL</li> </ul> </li> </ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2021-08-30  | <p>This issue is the thirtieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported real-time synchronization from MySQL to ElasticSearch.</li> <li>• Supported real-time synchronization from GaussDB primary/standby to GaussDB primary/standby.</li> <li>• Supported real-time synchronization for PostgreSQL 13.</li> <li>• Supported JSON-C message format in MySQL to Kafka synchronization scenario.</li> <li>• Supported exporting data comparison results.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• The following scenarios are in the open beta test phase. <ul style="list-style-type: none"> <li>- MySQL -&gt; GaussDB distributed</li> <li>- Oracle -&gt; GaussDB distributed</li> <li>- GaussDB distributed -&gt; Oracle</li> <li>- GaussDB distributed -&gt; GaussDB(DWS)</li> <li>- GaussDB distributed -&gt; Kafka</li> <li>- GaussDB distributed -&gt; GaussDB distributed</li> </ul> </li> <li>• Supported selecting DRS task exception notifications from the SMN topic.</li> </ul> |
| 2021-08-03  | <p>This issue is the twenty-ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported real-time synchronization from GaussDB(for MySQL) to Kafka.</li> <li>• Supported real-time synchronization from DDM to MySQL.</li> <li>• Supported real-time synchronization from PostgreSQL to GaussDB distributed in the cloud.</li> <li>• Supported real-time synchronization from PostgreSQL to GaussDB primary/standby in the cloud.</li> <li>• Supported filtering out the DELETE operation during out-of-cloud synchronization from MySQL to Kafka.</li> <li>• Supported column processing for real-time synchronization from MySQL to MySQL.</li> <li>• Supported user and permission synchronization and user comparison during PostgreSQL to PostgreSQL real-time synchronization.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• Supported commercial use of Oracle to GaussDB(DWS) synchronization.</li> </ul>                                                                                             |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2021-07-05  | <p>This issue is the twenty-eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported out-of-cloud synchronization from DDM to Kafka.</li> <li>• Supported flow control for MySQL to GaussDB(DWS) synchronization tasks.</li> <li>• Added permissions, allowing users to perform all operations except deleting DB instances.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• Charged fees for data transmission through a public network.</li> <li>• Adjusted fees for migrating databases out of the cloud.</li> <li>• Supported the commercial use of the following scenarios: <ul style="list-style-type: none"> <li>– Real-time synchronization from DDM to RDS for MySQL</li> <li>– Real-time synchronization from MySQL to RDS for PostgreSQL</li> <li>– Real-time synchronization from MySQL to GaussDB(for MySQL)</li> </ul> </li> </ul> |
| 2021-05-31  | <p>This issue is the twenty-seventeenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported tables without primary keys synchronized from MySQL to GaussDB(DWS).</li> <li>• Supported synchronization of some DDL operations during incremental synchronization from MySQL to GaussDB(DWS).</li> <li>• Supported synchronization of some DDL operations during incremental synchronization from PostgreSQL to PostgreSQL.</li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• Supported selecting DRS task exception notifications from the SMN topic.</li> </ul>                                                                                                                                                                                                                                                                                      |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2021-04-30  | <p>This issue is the twenty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported the following real-time synchronization scenarios: <ul style="list-style-type: none"> <li>- From Oracle to DDM.</li> <li>- From DDM to GaussDB(DWS)</li> <li>- From GaussDB distributed to GaussDB(DWS)</li> </ul> </li> </ul> <p>Changed the following content:</p> <ul style="list-style-type: none"> <li>• Supported the commercial use of the following scenarios: <ul style="list-style-type: none"> <li>- Real-time synchronization from MySQL to GaussDB(for MySQL)</li> <li>- Real-time synchronization from Oracle to GaussDB(for MySQL)</li> <li>- Real-time synchronization from MySQL to GaussDB(DWS)</li> <li>- Real-time synchronization from Oracle to RDS for MySQL</li> </ul> </li> </ul> |
| 2021-03-30  | <p>This issue is the twenty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported real-time synchronization from Oracle to GaussDB distributed.</li> <li>• Supported real-time synchronization from MySQL to GaussDB(for MySQL).</li> <li>• Supported real-time synchronization from PostgreSQL to GaussDB(DWS).</li> <li>• Supported incremental synchronization of MySQL databases.</li> <li>• Supported the pausing and resetting some synchronization task.</li> </ul>                                                                                                                                                                                                                                                                                                                   |
| 2021-01-30  | <p>This issue is the twenty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• DRS allows you to import, edit, and process files that are synchronized from Oracle to GaussDB(DWS) in real time.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 2020-12-30  | <p>This issue is the twenty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• DRS supports real-time synchronization from MySQL to GaussDB(DWS).</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| 2020-11-30  | <p>This issue is the twenty-second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported object editing and mapping during MySQL to GaussDB synchronization.</li> <li>• Supported searching objects when the user selects objects.</li> <li>• Supported setting the number of days after which an abnormal task can be automatically stopped.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                           |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2020-10-31  | <p>This issue is the twenty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Added the delay description in the DRS synchronization scenario.</li> </ul>                                                                                                                                                                                                      |
| 2020-09-30  | <p>This issue is the twentieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Optimized the minimum permissions of the destination database in the MySQL to GaussDB synchronization scenario.</li> <li>• Supported incremental startup of the GaussDB to MySQL synchronization task.</li> </ul>                                                                   |
| 2020-08-31  | <p>This issue is the nineteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported skipping data and pausing the synchronization from DDM to Oracle.</li> <li>• Supported configuration of the subnet for the synchronization instance.</li> </ul>                                                                                                          |
| 2020-07-31  | <p>This issue is the eighteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported MySQL multi-table row filtering during synchronization.</li> <li>• Allowed different users under the same tenant to manage their own DRS tasks, and the tasks are invisible to each other.</li> </ul>                                                                    |
| 2020-04-30  | <p>This issue is the seventeenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported exception diagnosis for heterogeneous synchronization tasks.</li> <li>• Supported resumable uploads by creating a synchronization task during Oracle to Kafka synchronization.</li> </ul>                                                                               |
| 2020-03-31  | <p>This issue is the sixteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported DDM to Oracle synchronization for the first time.</li> <li>• Supported MySQL to PostgreSQL synchronization over public networks.</li> <li>• Provided the task pausing function.</li> <li>• Supported MySQL to GaussDB(DWS) synchronization for the first time.</li> </ul> |

| Released On | Description                                                                                                                                                                                                                                                                                                                                  |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2020-02-29  | <p>This issue is the fifteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported pagination and querying for data synchronization.</li> <li>• Supported the change of the flow control mode after the task is started.</li> <li>• Supported resetting passwords.</li> </ul> |
| 2020-01-30  | <p>This issue is the fourteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported MySQL 8.0 for the first time.</li> <li>• Supported forcing tasks to stop.</li> </ul>                                                                                                      |
| 2019-12-30  | <p>This issue is the thirteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported synchronization from self-built MySQL databases to self-built MySQL databases.</li> </ul>                                                                                                 |
| 2019-11-30  | <p>This issue is the twelfth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported synchronizing MySQL out of the cloud.</li> </ul>                                                                                                                                             |
| 2019-10-30  | <p>This issue is the eleventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported tag management.</li> </ul>                                                                                                                                                                  |
| 2019-09-30  | <p>This issue is the tenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported compute resources selection in the value comparison function of MySQL.</li> <li>• Charged users for creating MySQL synchronization tasks.</li> </ul>                                           |
| 2019-08-30  | <p>This issue is the ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported the overwrite policy when conflicts occur during MySQL synchronization.</li> <li>• Supported controlling the database replication rate in the VPN and dedicate connect scenarios.</li> </ul>   |
| 2019-07-30  | <p>This issue is the eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Canceled the retry function and supported resetting and resuming MySQL synchronization tasks.</li> </ul>                                                                                                |
| 2019-06-30  | <p>This issue is the seventh official release, which incorporates the following change:</p> <ul style="list-style-type: none"> <li>• Supported resetting MySQL synchronization tasks.</li> </ul>                                                                                                                                             |

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                              |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2019-04-30  | <p>This issue is the sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported VPN and Dedicated Connect networks in synchronization scenarios.</li> </ul>                                                                                                                                                                                                |
| 2019-02-28  | <p>This issue is the fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported filtering DROP DATABASE in MySQL to MySQL synchronization.</li> <li>• Displayed mapping information in the synchronization scenario.</li> </ul>                                                                                                                            |
| 2019-01-19  | <p>This issue is the fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported editing the MySQL to MySQL synchronization task.</li> <li>• Supported MySQL to PostgreSQL synchronization.</li> <li>• Visualized data conflicts in data synchronization scenarios.</li> <li>• Supported table mapping for MySQL to PostgreSQL synchronization.</li> </ul> |
| 2018-12-30  | <p>This issue is the third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Optimized the synchronization mode and object.</li> </ul>                                                                                                                                                                                                                            |
| 2018-11-30  | <p>This issue is the second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"> <li>• Supported data synchronization.</li> <li>• Classified check items.</li> <li>• Supported deleting tasks in batches.</li> </ul>                                                                                                                                                       |
| 2018-10-31  | <p>This issue is the first official release.</p>                                                                                                                                                                                                                                                                                                                                                                         |