

Data Replication Service

Real-Time Migration

Issue 49
Date 2024-11-30



Copyright © Huawei Cloud Computing 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 Cloud Computing Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property 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 Cloud 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 Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 Migration Overview.....	1
2 To the Cloud.....	6
2.1 From MySQL to MySQL.....	6
2.2 From MySQL to DDM.....	29
2.3 From MySQL to GaussDB(for MySQL).....	44
2.4 From MySQL Schema and Logic Table to DDM.....	64
2.5 From MongoDB to DDS.....	78
2.6 From MongoDB to GeminiDB Mongo.....	100
2.7 From Redis to GeminiDB Redis.....	116
2.8 From Redis Cluster to GeminiDB Redis.....	129
3 Out of the Cloud.....	142
3.1 From MySQL to MySQL.....	142
3.2 From DDS to MongoDB.....	163
3.3 From GeminiDB Redis to Redis.....	180
3.4 From GeminiDB Redis to Redis Cluster.....	192
4 Task Management.....	205
4.1 Creating a Migration Task.....	205
4.2 Querying the Migration Progress.....	223
4.3 Viewing Migration Logs.....	226
4.4 Data Comparison (Comparing Migration Items).....	227
4.5 Managing Objects.....	238
4.5.1 Migrating Accounts.....	238
4.5.2 Parameters for Comparison.....	241
4.5.3 Viewing Migration Objects.....	245
4.6 Task Life Cycle.....	246
4.6.1 Viewing Task Details.....	246
4.6.2 Modifying Task Information.....	246
4.6.3 Modifying Connection Information.....	248
4.6.4 Modifying the Flow Control Mode.....	249
4.6.5 Editing a Migration Task.....	251
4.6.6 Resuming a Migration Task.....	256
4.6.7 Disabling Read-Only of the Destination Database.....	257

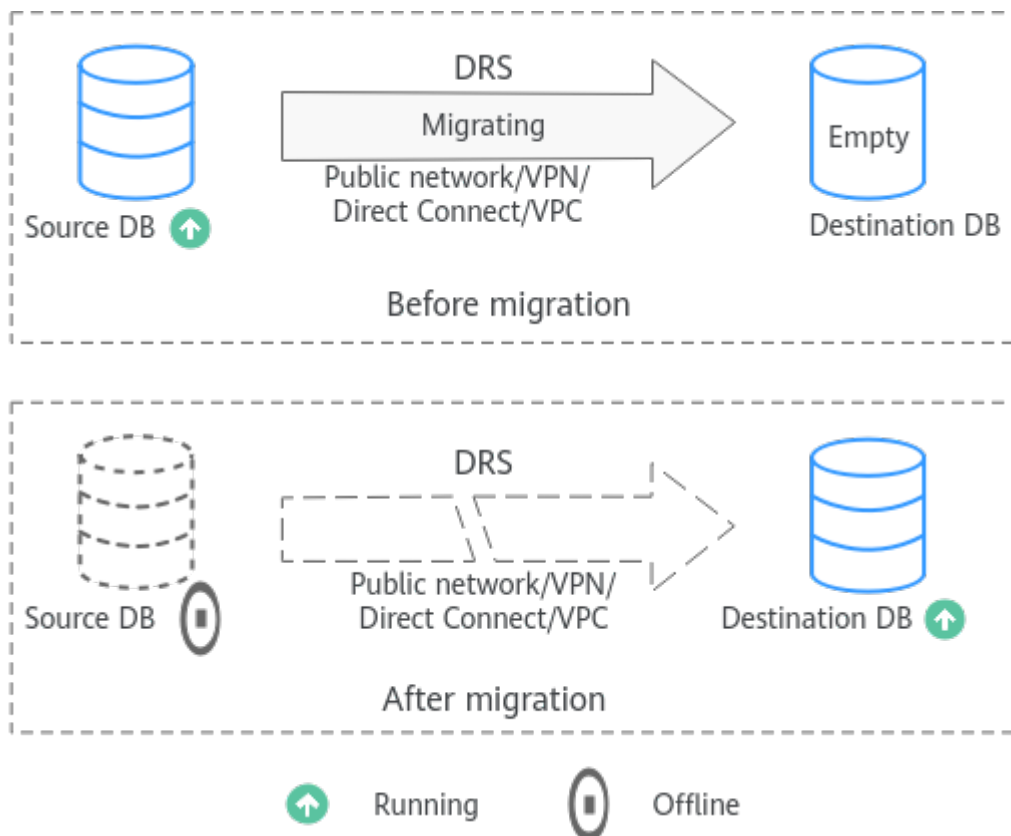
4.6.8 Resetting a Migration Task.....	258
4.6.9 Pausing a Migration Task.....	259
4.6.10 Cloning a Migration Task.....	261
4.6.11 Stopping a Migration Task.....	261
4.6.12 Deleting a Migration Task.....	263
4.6.13 Importing Tasks in Batches.....	264
4.6.14 Upgrading the Version of a DRS Task.....	265
4.6.15 Task Statuses.....	268
5 Tag Management.....	270
6 Connection Diagnosis.....	272
7 Interconnecting with CTS.....	274
7.1 Key Operations Recorded by CTS.....	274
7.2 Viewing Traces.....	274
8 Interconnecting with Cloud Eye.....	276
8.1 Supported Metrics.....	276
8.2 Configuring Alarm Rules.....	281
8.3 Viewing Monitoring Metrics.....	282
9 Interconnecting with LTS.....	284
9.1 Log Reporting.....	284
9.2 Viewing and Downloading Logs.....	285
10 Appendix.....	288
10.1 General Usage Suggestions.....	288
10.2 DRS Pre-Check Items.....	289

1 Migration Overview

With DRS, you can migrate data from sources to destinations in real time. You create a replication instance to connect to both the source and destination and configure objects to be migrated. DRS will help you compare metrics and data between source and destination, so you can determine the best time to switch to the destination database while minimizing service downtime.

DRS supports incremental migration, so you can replicate ongoing changes to keep sources and destinations in sync while minimizing the impact of service downtime and migration.

Figure 1-1 Real-time migration process



Supported Database Types

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

Table 1-1 Migration schemes

Source DB	Destination DB	Migration Type	Documentation
<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	Full Full+Incremental	From MySQL to MySQL (To the cloud)
	DDM	Full Full+Incremental	From MySQL to DDM
	GaussDB(for MySQL)	Full Full+Incremental	From MySQL to GaussDB(for MySQL)
<ul style="list-style-type: none"> RDS for MySQL 	<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds 	Full Full+Incremental	From MySQL to MySQL (Out of the cloud)
<ul style="list-style-type: none"> On-premises MongoDB databases MongoDB databases on an ECS MongoDB database on other clouds DDS 	DDS	Full Full+Incremental	From MongoDB to DDS
<ul style="list-style-type: none"> On-premises MongoDB databases MongoDB databases on an ECS MongoDB database on other clouds 	GeminiDB Mongo	Full Full+Incremental: <ul style="list-style-type: none"> Replica set -> Replica set Replica set -> Cluster Cluster -> Cluster 	From MongoDB to GeminiDB Mongo

Source DB	Destination DB	Migration Type	Documentation
<ul style="list-style-type: none"> • DDS 	<ul style="list-style-type: none"> • On-premises MongoDB databases • MongoDB databases on an ECS • MongoDB database on other clouds 	Full Full+Incremental	From DDS to MongoDB
<ul style="list-style-type: none"> • DDM 	DDM	Full Full+Incremental	From MySQL Schema and Logic Table to DDM
<ul style="list-style-type: none"> • On-premises official open-source single-node or master/standby Redis • ECS-hosted official open-source single-node or master/standby Redis 	GeminiDB Redis	Full Full+Incremental	From Redis to GeminiDB Redis
<ul style="list-style-type: none"> • On-premises open-source Redis Cluster • ECS-hosted open-source Redis Cluster 	GeminiDB Redis	Full+Incremental	From Redis Cluster to GeminiDB Redis

Source DB	Destination DB	Migration Type	Documentation
GeminiDB Redis	<ul style="list-style-type: none"> On-premises open-source Codis ECS-hosted open-source Codis On-premises single-node or master/standby Redis ECS-hosted single-node or master/standby Redis GeminiDB Redis 	Full+Incremental	From GeminiDB Redis to Redis
	<ul style="list-style-type: none"> On-premises Redis Cluster ECS-hosted Redis Cluster 	Full+Incremental	From GeminiDB Redis to Redis Cluster

Other Cloud Migration Solutions

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

Table 1-2 Migration solutions

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

Source Database Cloud Vendor	Source Database Type	Destination Database (Huawei Cloud)	Solution
Azure	Database for MySQL	RDS for MySQL	From MySQL to MySQL
Azure	Database for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Google	Cloud SQL for MySQL	RDS for MySQL	From MySQL to MySQL
Google	Cloud SQL for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Alibaba Cloud	ApsaraDB RDS for MySQL	RDS for MySQL	From MySQL to MySQL
Alibaba Cloud	ApsaraDB RDS for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Alibaba Cloud	PolarDB for MySQL	RDS for MySQL	From MySQL to MySQL
Alibaba Cloud	PolarDB for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Tencent Cloud	TencentDB for MySQL	RDS for MySQL	From MySQL to MySQL
Tencent Cloud	TencentDB for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)
Tencent Cloud	TDSQL-C for MySQL	RDS for MySQL	From MySQL to MySQL
Tencent Cloud	TDSQL-C for MySQL	GaussDB(for MySQL)	From MySQL to GaussDB(for MySQL)

2 To the Cloud

2.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 2-1 Supported databases

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

Database Account Permission Requirements

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

Table 2-2 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	<p>The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT</p> <p>If the source database version is 8.0, the user must have the SELECT permission for the mysql.user table. If the source database version is 5.7 or earlier, the user must have the SELECT permission for the MySQL system database. If the source database is a Alibaba Cloud database, the user must have the SELECT permission for both mysql.user and mysql.user_view.</p>	<p>The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <p>If the source database version is 8.0, the user must have the SELECT permission for the mysql.user table. If the source database version is 5.7 or earlier, the user must have the SELECT permission for the MySQL system database. If the source database is a Alibaba Cloud database, the user must have the SELECT permission for both mysql.user and mysql.user_view. If the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup.</p>
Destination database user	<p>The user must have the following minimum permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</p> <p>To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database.</p>	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- [Table 2-2](#) lists the minimum permissions required by a DRS task. If you need to migrate the grant permission through a DRS task, ensure that the connection account of the DRS task has the corresponding permission. Otherwise, the destination database user may not be authorized due to grant execution failure. For example, if the connection account of the DRS task does not require the process permission, but you need to migrate the process permission through a DRS task, ensure that the connection account of the DRS task has the process permission.

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-3](#). DRS will automatically check the objects you selected before the migration.

Table 2-3 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: <code>.'<>^/</code> • The source table and view names cannot contain non-ASCII characters, or the following characters: <code>'<>^/</code> • You can select table-, database-, or instance-level (all-level) migration. • Databases, tables, users, views, indexes, constraints, functions, stored procedures, triggers, and events can be migrated. • Only MyISAM and InnoDB tables can be migrated. • The system database and event statuses cannot be migrated. • Non-standard floating-point data that can be written in loose mode but cannot be written in strict mode cannot be migrated. Such non-standard floating-point data may fail to be hit, causing data migration failures. • Full and incremental migrations do not support invisible columns. Invisible columns can be migrated 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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 2-4](#).

Table 2-4 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database cannot be RDS read replicas of Alibaba Cloud.• The source database names cannot contain non-ASCII characters, or the following characters: '<code><>\</code>'• The source table and view names cannot contain non-ASCII characters, or the following characters: '<code><>\</code>'• If the source database is an on-premises database and has Percona Server for MySQL 5.6.x or Percona Server for MySQL 5.7.x installed, the memory manager must use Jemalloc to prevent Out of Memory errors caused by frequent queries on system tables.• The source database does not support the mysql binlog dump command.• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.• 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.• Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> ● During the migration, 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. ● During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. ● During migration, do not write the statement-based binlog into the source database. ● During migration, do not clear the binlog in the source database. ● During migration, do not create a database named ib_logfile in the source database. ● During migration, do not change the floating IP address of the primary node of the destination database to that of the read replica. If so, data will be written to the read replica. As a result, data in the destination database is inconsistent with that in the source database. ● During the migration, 0 cannot be written to the auto-increment primary key column in the source database. Otherwise, the data of the auto-increment column in the source database is inconsistent with that in the destination database. ● During the migration, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the migration task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a migration task again if the source MySQL database is upgraded across major versions. ● If index migration 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 migration</p> <ul style="list-style-type: none"> ● During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> ● During an incremental migration of table-level objects, renaming tables is not supported. ● During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database.

Type	Restrictions
	<ul style="list-style-type: none"> • During an incremental migration, 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. • DDL statements are supported in the incremental migration phase. <p>Stopping a task</p> <ul style="list-style-type: none"> • Stop a task normally. The selected events and triggers are migrated while the migration task proceeds to the final stage. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the migration status reported by the migration log. • Forcibly stop a task. If you forcibly stop a task, DRS resources will be released and triggers and events will not be migrated. You need to manually migrate triggers and events. If you want DRS to migrate triggers and events, restore the DRS task first. After the task status becomes normal, stop the task. <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 migration 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. ● When creating multiple migration tasks in the many-to-one scenario, ensure that the read and write settings of the destination database are consistent in these tasks. ● 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 a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The destination database cannot be restored to a point in time when a full migration was being performed. ● If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem. ● If the source 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. ● The destination database of a migration task can be set to Read-only or Read/Write. <ul style="list-style-type: none"> – Read-only: During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration. – Read/Write: During the migration, the destination instance can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may

Type	Restrictions
	result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks.

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 migration, 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](#).
- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

This section uses the migration from MySQL to RDS for MySQL in a VPC as an example to describe how to configure a migration task on the DRS console.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-1 Migration task information

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

Region

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

Project

* Task Name ⓘ

Description
0/256 ⓘ

Table 2-5 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-2 Replication instance information

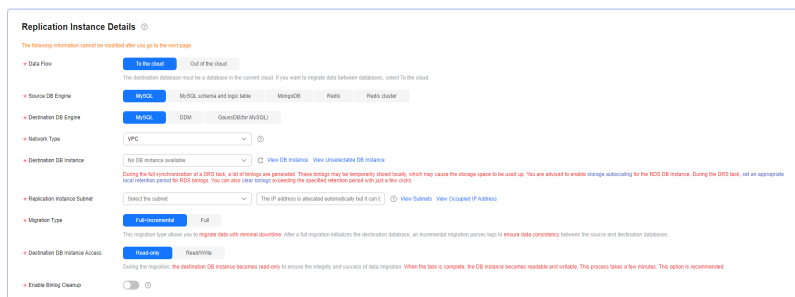


Table 2-6 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination DB is on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	Select VPC Network. Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network . <ul style="list-style-type: none"> – VPC is suitable for migrations between cloud databases of the same account in the same region and VPC. – Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases. – VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection. For details about networks, see Preparations .

Parameter	Description
Destination DB Instance	The RDS DB instance you created.
Replication Instance Subnet	<p>The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Destination DB Instance Access	<ul style="list-style-type: none"> <li data-bbox="691 719 1441 1424"> <p>– Read-only</p> <p>During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.</p> <p>If a migration task fails, the destination database restores to the read/write state. After the migration task is resumed, the destination database becomes read-only.</p> <p>The access settings of the destination instance are also affected by the access settings of the DB instance itself. Therefore, you are advised not to set the access settings of the DB instance on the RDS console.</p> <p>If the destination database is set to read-only for a migration task, you can disable read-only on the task details page. After read-only is disabled, the destination database cannot be set back to read-only. Even if the migration task fails and then resumed and started, the destination database cannot be set to read-only.</p> <li data-bbox="691 1435 1441 1872"> <p>– Read/Write</p> <p>During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks. Set the destination database to Read/Write only when you need to modify other data in the database during the migration.</p> <p>The task cannot be modified after being created.</p>

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects and data in non-system databases are migrated to the destination database at one time. The objects include tables, views, and stored procedures. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database. - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full migration and disabled during the incremental migration.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 2-3 AZ




Table 2-7 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 2-4 Enterprise projects and tags

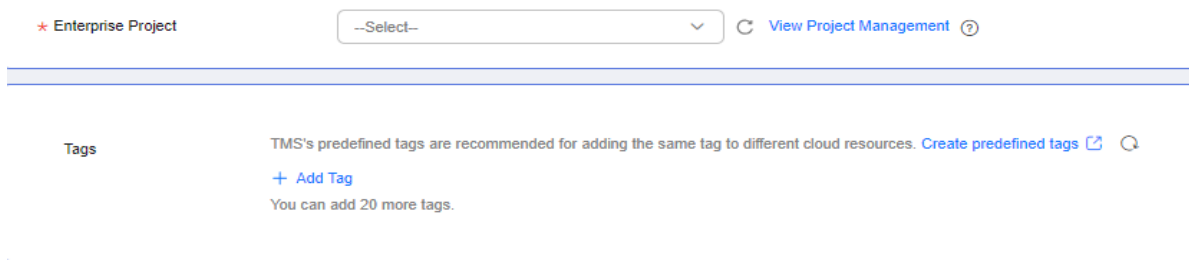


Table 2-8 Enterprise Project and Tags

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 projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 2-5 DNS Server

Configure Your Own DNS Server 

DNS Server

DNS Server IP Address

Table 2-9 DNS server information

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

 **NOTE**

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

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

- Scenario 1: Self-built database - source database configuration

Figure 2-6 Self-built database - 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 Username

Database Password

SSL Connection

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

Table 2-10 Self-built database - source database information

Parameter	Description
Database Type	Select Self-built database .
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 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 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 2-7 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.

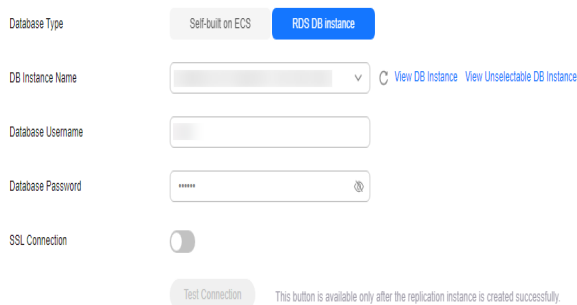


Table 2-11 RDS DB instance - source database information

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the RDS DB instance to be migrated as the source DB instance.
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.

- Destination database configuration

Figure 2-8 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Migrate Definer to User Yes No

SSL Connection

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

Table 2-12 Destination database settings

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

Parameter	Description
Database Password	The password for the database username.
Migrate Definer to User	<p>Indicates whether to migrate the Definers of all source database objects to the destination database user entered during the connection test.</p> <ul style="list-style-type: none"> - Yes The Definers of all source database objects will be migrated to the user. Other users do not have permissions for database objects unless these users are authorized. For details about authorization, see How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration? <p>For example, if the view is <code>CREATE ALGORITHM=UNDEFINED DEFINER='username'@'%` SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> before migration,</p> <p>it is converted to <code>CREATE ALGORITHM=UNDEFINED DEFINER='drsUser'@'%` SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> after the migration.</p> <p>drsUser indicates the destination database user used for testing the connection.</p> <ul style="list-style-type: none"> - No The Definers of all source database objects will not be changed. You need to migrate all accounts and permissions of the source database in the next step. Note that if the Definer account is not found in the destination database, unavailable objects will be created. <p>For details about Definer, see the MySQL official document.</p>
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 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 Task** page, select the accounts and objects to be migrated, and click **Next**.

Figure 2-9 Migration type

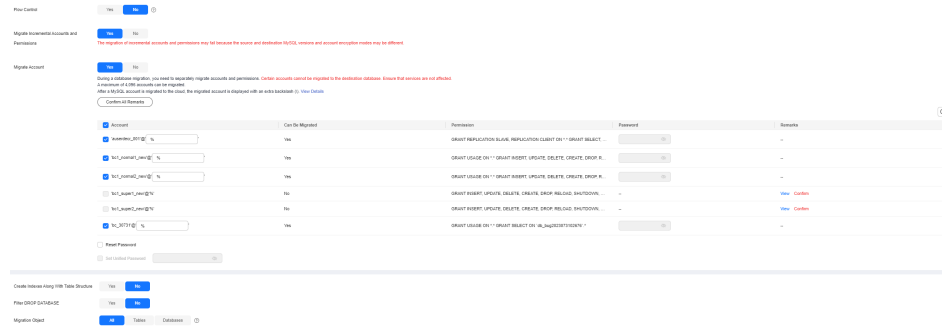
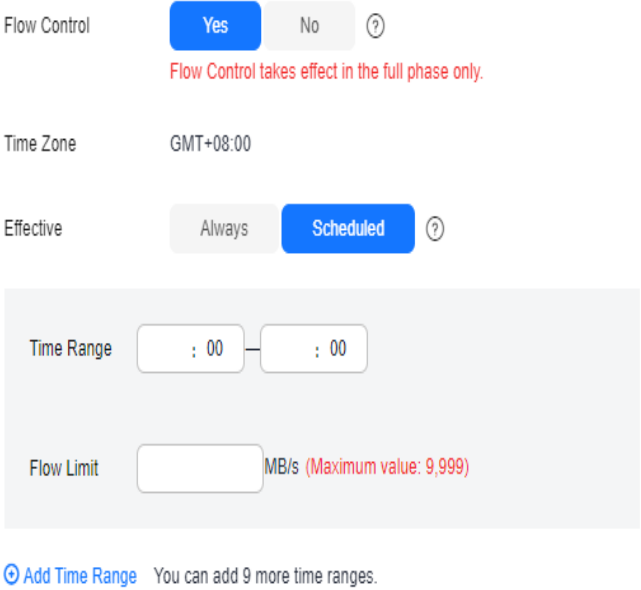



Table 2-13 Migration types and objects

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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 2-10 Configuring flow control</p>  <p>No The migration 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>

Parameter	Description
	<p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Migrate Incremental Accounts and Permissions	<p>Indicates whether to migrate incremental accounts and permissions during database migration.</p> <ul style="list-style-type: none">● Yes All incremental accounts and permissions will be migrated. The migration of incremental accounts and permissions may fail because the source and destination database versions and account encryption modes may be different.● No All incremental accounts and permissions will be filtered out during the migration.
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none">● Yes If you need to migrate accounts, see Migrating Accounts.● No During migration, accounts, permissions, and passwords are not migrated.
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full migration phase.</p> <ul style="list-style-type: none">● Yes: Indexes are migrated when the table structure is migrated in the full migration phase.● No: Indexes are migrated separately after data migration.
Filter DROP DATABASE	<p>To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations.</p> <ul style="list-style-type: none">● If you select Yes, any database deletion operations performed on the source database are not migrated during data migration.● If you select No, related operations are migrated to the destination database during data migration.

Parameter	Description
Migration Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> ● All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. ● Tables: The selected table-level objects will be migrated. ● Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. ● 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 migration object cannot contain spaces. ● To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Step 6 Compare source and destination parameters.

By comparing common and performance parameters for the source databases against those of the destination databases, you can help ensure that services will not change after a migration is completed. You can determine whether to use this function based on service requirements. It mainly ensures that services are not affected after a migration is completed.

- This process is optional, so you can click **Next** to skip the comparison.
- Compare common parameters:
If the common parameter values in the comparison results are inconsistent, click **Save Change** to change the destination database values to be the same as those of the source database.

Figure 2-11 Modifying common parameters

Parameter Name	Source Database Value	Destination Database Value	Result
connect_timeout	10	10	Consistent
explicit_defaults_for_timestamp	OFF	OFF	Consistent
innodb_flush_log_at_trx_commit	1	1	Consistent
innodb_lock_wait_timeout	50	50	Consistent
innodb_read_io_threads	2500	153	Inconsistent
innodb_read_timeout	30	30	Consistent
innodb_write_timeout	60	60	Consistent
transaction_isolation	REPEATABLE-READ	REPEATABLE-READ	Consistent

Performance parameter values in both the source and destination databases can be the same or different.

- If you need to change the performance parameter values that are consistent in the comparison results to different values, locate the target parameter, enter values in the **Change To** column, and click **Save Change** in the upper left corner.
- If you want to make the performance parameter values of the source and destination database be the same:
 - Click **Use Source Database Value**.
DRS automatically makes the destination database values the same as those of the source database.

Figure 2-12 One-click modification

Parameter Name	Source Database Value	Destination Database Value	Change To	Allowed Destination Database Value	Result	
innodb_io_capacity	32768	32768	8	4096 - 32768	4096-1077216	Consistent
innodb_io_capacity_max	32768	32768	8	4096 - 32768	4096-1077216	Consistent
innodb_read_io_threads	8388008	8388008		0-18467440733855615	0-18467440733855615	Consistent
innodb_write_io_threads	2	2		1-64	1-64	Consistent
innodb_buffer_pool_size	429487296	429487296	16	25602896 - 429487296	1073741024-4471847873	Consistent
innodb_buffer_pool_instances	1	1	8	8	8	Consistent
innodb_read_ahead_threads	202144	202144	8K	4096 - 202144	8192-214791952	Consistent
innodb_read_ahead_threads_max	524288	524288	128	4096 - 524288	4096-214791952	Consistent
innodb_read_ahead_buffer_size	202144	202144			32768-18467440733855615	Consistent
innodb_log_block_size	1	1			0-429487296	Consistent

NOTE

You can also manually enter parameter values.

- Click **Save Change** to save your changes.
The system changes the parameter values based on your settings for the destination database values. After the modification, the list is updated automatically.

Figure 2-13 Performance parameters

Parameter Name	Source Database Value	Destination Database Value	Change To	Allowed Destination Database V.	Result
<input type="checkbox"/> innodb_log_file_size	32768	32768	0	4096 + 32768	Consistent
<input type="checkbox"/> innodb_log_file_no	32768	32768	0	4096 + 32768	Consistent
<input type="checkbox"/> innodb_log_file_no	8388496	8388496		0-1644814487370551615	Consistent
<input type="checkbox"/> innodb_log_file_instances	2	2		1-64	Consistent
<input checked="" type="checkbox"/> innodb_log_file_size	429497296	429497296	0	268435456 + 2147483648	Consistent
<input type="checkbox"/> innodb_log_file_size	1300000	1300000		0-51-3000	Consistent
<input type="checkbox"/> innodb_log_file_size	262144	262144	0	4096 + 262144	Consistent
<input type="checkbox"/> innodb_log_file_size	524288	524288	0	4096 + 524288	Consistent
<input type="checkbox"/> innodb_log_file_size	262144	262144		32768-1644814487370551615	Consistent
<input type="checkbox"/> innodb_log_file_size	1	1		0-429497296	Consistent

Some parameters in the destination database require a restart before the changes can take effect. The system will display these as being inconsistent. In addition, restart the destination database before the migration task is started or after the migration task is completed. To minimize the impact of this restart on your services, it is recommended that you schedule a specific time to restart the destination database after the migration is complete.

For details about how to set parameters during a comparison, see [Parameters for Comparison](#).

iii. Click **Next**.

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

Figure 2-14 Task startup settings

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

Apr 11, 2024 03:00:00 ⓘ GMT+08:00

In the full migration phase, the DRS task occupies performance resources. To ensure services and tasks can run normally, you are advised to start the task during off-peak hours and reserve two to three days for data verification.

Send Notifications: ⓘ

* SMN Topic: ⓘ ⓘ

Delay Threshold (s): ⓘ


Data Exception Notification:

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

Table 2-14 Task startup settings

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

Step 8 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.2 From MySQL to DDM

Supported Source and Destination Databases

Table 2-15 Supported databases

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

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-16](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 2-16 Database account permission

Type	Full Migration	Full+Incremental Migration
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. If the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup.
Destination database user	<ul style="list-style-type: none">• 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 migrated.	

NOTE

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

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-17](#). DRS will automatically check the objects you selected before the migration.

Table 2-17 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none"> ● The source database names cannot contain non-ASCII characters, or the following characters: . '<' '>' '^' ● The column names in the source database tables cannot end with a backslash (\). ● The source table and view names cannot contain non-ASCII characters, or the following characters: '<' '>' '^' ● You can select table-level migration. ● Only the source database data can be migrated. The table structure and other database objects of the source database cannot be migrated. ● Only MyISAM and InnoDB tables can be migrated. ● The system database and event statuses cannot be migrated. ● Tables without primary keys cannot be migrated. ● Non-standard floating-point data that can be written in loose mode but cannot be written in strict mode cannot be migrated. Such non-standard floating-point data may fail to be hit, causing data migration failures. ● If the source database is MySQL 8.0, tables with function indexes cannot be migrated. ● Full and incremental migrations do not support invisible columns. Invisible columns can be migrated 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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 2-18](#).

Table 2-18 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: `.'<'>\/` • The source table and view names cannot contain non-ASCII characters, or the following characters: `<'>\/` • If the source database is an on-premises database and has Percona Server for MySQL 5.6.x or Percona Server for MySQL 5.7.x installed, the memory manager must use Jemalloc to prevent Out of Memory errors caused by frequent queries on system tables. • The source database does not support the mysql binlog dump command. • The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency. • Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. • 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. • Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • During the migration, do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. • During migration, do not write the statement-based binlog into the source database. • During migration, do not clear the binlog in the source database. • During migration, do not create a database named ib_logfile in the source database. • During the migration, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the migration task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a migration task again if the source MySQL database is upgraded across major versions. <p>Full migration</p> <ul style="list-style-type: none"> • During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> • During an incremental migration of table-level objects, renaming tables is not supported. • During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database. • During an incremental migration, 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. • During an incremental migration, the following DDL statements are supported: ADD COLUMN, MODIFY COLUMN, CHANGE COLUMN, CREATE INDEX, DROP INDEX, RENAME INDEX and ADD INDEX. <p>Stopping a task</p> <ul style="list-style-type: none"> • 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.

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 migration task cannot be created. • When creating multiple migration tasks in the many-to-one scenario, ensure that the read and write settings of the destination database are consistent in these tasks. • 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 a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. • If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. • 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. • Ensure that the destination database is empty before starting the migration. Otherwise, data in the destination may be overwritten during incremental migration. • If the destination DB instance uses columns of the TIMESTAMP or DATETIME data type as its sharding key, the seconds precision of the column is removed after the migration.

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 migration, 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](#).
- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-15 Migration task information

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

Region: [Dropdown menu] .. v

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

Project: [Dropdown menu] .. v

* Task Name: [Text input: DRS-5678] ⓘ

Description: [Text area] ⓘ

0/256

Table 2-19 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-16 Replication instance information

Replication Instance Details ⓘ

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

• Data Floor: **To the cloud** (Selected) | Out of the cloud

The destination database must be a database in the current cloud. If you want to migrate data between databases, select to the cloud.

• Source DB Engine: **MySQL** | MySQL schema and topic table | MongoDB | Redis | Redis cluster

• Destination DB Engine: MySQL | **Click** | Click to edit MySQL

• Network Type: Public network

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

• Destination DB Instance: No DB instance available | [View DB Instance](#) | [View Unavailable DB Instance](#)

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

• Migration Type: **Full** | [Full](#) | [Full-Incremental](#) | [Full](#)

This migration type allows you to migrate data with minimal downtime. After a full migration replicates the destination database, an incremental migration passes logs to ensure data consistency between the source and destination databases.

• Security EP: [Dropdown menu] | [Create an EP](#)

Table 2-20 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination is a DB instance on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select DDM .
Network Type	Available options: VPC , Public network , and VPN or Direct Connect . By default, the value is Public network . <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region.- VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions.- Public network is suitable for migration from on-premises databases or external cloud databases to destination databases.
Destination DB Instance	The DDM instance you created.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects in non-system databases are migrated to the destination database at one time, including tables, views, stored procedures, and triggers. <p>NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.</p> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 2-17 AZ

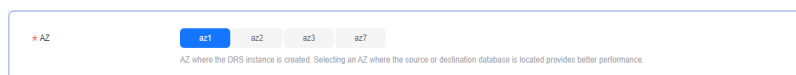


Table 2-21 Task AZ

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 2-18 Enterprise projects and tags




Table 2-22 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 projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database configuration

Figure 2-19 Source database information

Source Database

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

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Table 2-23 Source database settings

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

 **NOTE**

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

- Destination database configuration

Figure 2-20 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

✔ Test successful

Table 2-24 Destination database settings

Parameter	Description
DB Instance Name	The DDM instance selected when you create the migration task. This parameter cannot be changed.
Database Username	The username for accessing the destination DDM 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 Task** page, select migration objects and click **Next**.

Figure 2-21 Migration object

Flow Control ⓘ

Migration Object ⓘ

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 ⓘ

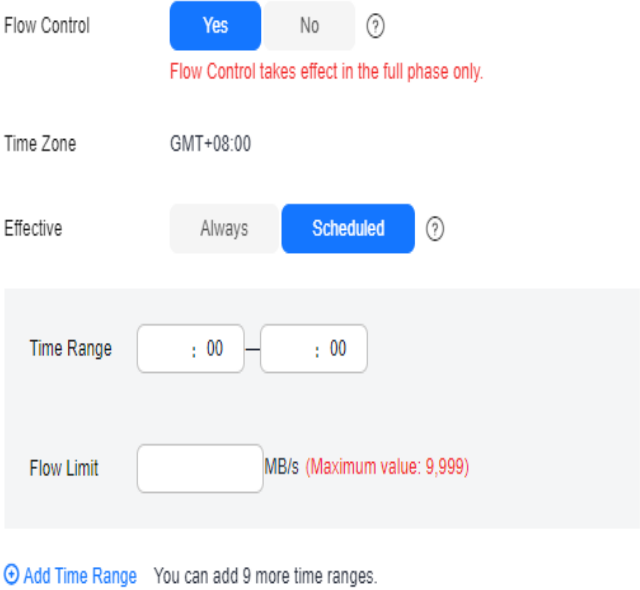
Search the expanded database using regular expressions. 🔍


Select All ⓘ

Search the expanded database using regular expressions. 🔍

- db_tatus_check_009 database

Table 2-25 Migration 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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 2-22 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> - Flow control mode takes effect only during a full migration. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can migrate table-level objects to the destination database based on service requirements.</p> <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If an object name contains spaces, the spaces before and after the object name are not displayed. If there are two or more consecutive spaces in the middle of the object name, only one space is displayed. • The name of the selected migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

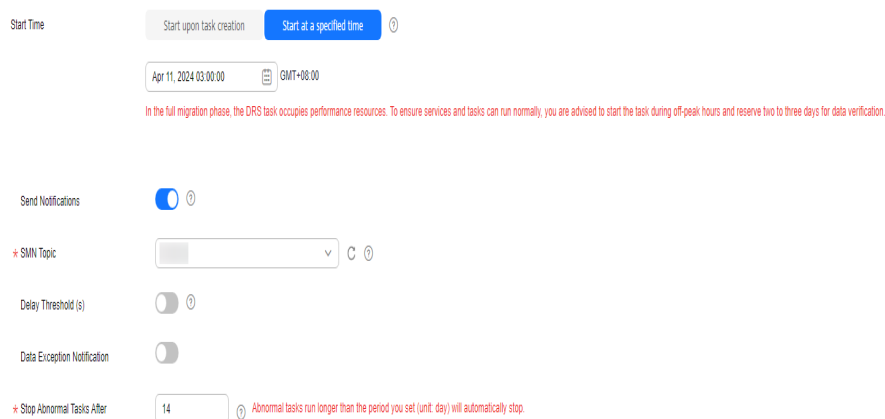
- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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

Figure 2-23 Task startup settings



Start Time

Start upon task creation Start at a specified time ⓘ

Apr 11, 2024 03:00:00 ⓘ GMT+08:00

In the full migration phase, the DRS task occupies performance resources. To ensure services and tasks can run normally, you are advised to start the task during off-peak hours and reserve two to three days for data verification.

Send Notifications ⓘ

* SMN Topic ⓘ

Delay Threshold (s) ⓘ


Data Exception Notification

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

Table 2-26 Task startup settings

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.3 From MySQL to GaussDB(for MySQL)

Supported Source and Destination Databases

Table 2-27 Supported databases

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

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-28](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 2-28 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	<p>The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT</p> <p>If the source database version is 8.0, the user must have the SELECT permission for the mysql.user table. If the source database version is 5.7 or earlier, the user must have the SELECT permission for the MySQL system database. If the source database is a Alibaba Cloud database, the user must have the SELECT permission for both mysql.user and mysql.user_view.</p>	<p>The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <p>If the source database version is 8.0, the user must have the SELECT permission for the mysql.user table. If the source database version is 5.7 or earlier, the user must have the SELECT permission for the MySQL system database. If the source database is a Alibaba Cloud database, the user must have the SELECT permission for both mysql.user and mysql.user_view. If the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup.</p>
Destination database user	<p>The user must have the following minimum permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</p> <p>To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database.</p>	

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- [Table 2-28](#) lists the minimum permissions required by a DRS task. If you need to migrate the grant permission through a DRS task, ensure that the connection account of the DRS task has the corresponding permission. Otherwise, the destination database user may not be authorized due to grant execution failure. For example, if the connection account of the DRS task does not require the process permission, but you need to migrate the process permission through a DRS task, ensure that the connection account of the DRS task has the process permission.

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-29](#). DRS will automatically check the objects you selected before the migration.

Table 2-29 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none"> • The source database names cannot contain non-ASCII characters, or the following characters: <code>.'<>^\"</code> • The source table and view names cannot contain non-ASCII characters, or the following characters: <code>'<>^\"</code> • You can select table-, database-, or instance-level (all-level) migration. • Databases, tables, users, views, indexes, constraints, functions, stored procedures, triggers, and events can be migrated. • Only MyISAM and InnoDB tables can be migrated. • The system database and event statuses cannot be migrated. • Non-standard floating-point data that can be written in loose mode but cannot be written in strict mode cannot be migrated. Such non-standard floating-point data may fail to be hit, causing data migration failures. • Full and incremental migrations do not support invisible columns. Invisible columns can be migrated 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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 2-30](#).

Table 2-30 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database cannot be RDS read replicas of Alibaba Cloud.• The source database names cannot contain non-ASCII characters, or the following characters: '<code><</code>'>'\'• The source table and view names cannot contain non-ASCII characters, or the following characters: '<code><</code>'>'\'• The column names in the source database tables cannot end with a backslash (\).• If the source database is an on-premises database and has Percona Server for MySQL 5.6.x or Percona Server for MySQL 5.7.x installed, the memory manager must use Jemalloc to prevent Out of Memory errors caused by frequent queries on system tables.• The source database does not support the mysql binlog dump command.• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.• 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.• Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • During the migration, 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. • During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. • During migration, do not write the statement-based binlog into the source database. • During migration, do not clear the binlog in the source database. • During migration, do not create a database named ib_logfile in the source database. • During migration, do not change the floating IP address of the primary node of the destination database to that of the read replica. If so, data will be written to the read replica. As a result, data in the destination database is inconsistent with that in the source database. • During the migration, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the migration task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a migration task again if the source MySQL database is upgraded across major versions. • If index migration 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 migration</p> <ul style="list-style-type: none"> • During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> • During an incremental migration of table-level objects, renaming tables is not supported. • During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database. • During an incremental migration, 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.

Type	Restrictions
	<ul style="list-style-type: none">• DDL statements are supported in the incremental migration phase. <p>Stopping a task</p> <ul style="list-style-type: none">• Stop a task normally. The selected events and triggers are migrated while the migration task proceeds to the final stage. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the migration status reported by the migration log.• Forcibly stop a task. If you forcibly stop a task, DRS resources will be released and triggers and events will not be migrated. You need to manually migrate triggers and events. If you want DRS to migrate triggers and events, restore the DRS task first. After the task status becomes normal, stop the task. <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 migration task cannot be created. ● The destination database must be a primary/standby GaussDB(for MySQL) instance. ● When creating multiple migration tasks in the many-to-one scenario, ensure that the read and write settings of the destination database are consistent in these tasks. ● 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 a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent. ● The destination database cannot be restored to a point in time when a full migration was being performed. ● If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem. ● If the source 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. ● The destination database of a migration task can be set to Read-only or Read/Write. <ul style="list-style-type: none"> - Read-only: During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration. - Read/Write: During the migration, the destination instance can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks.

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 migration, 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](#).
- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-24 Migration task information

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

Region

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

Project

* Task Name

Description

0/256

Table 2-31 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-25 Replication instance information

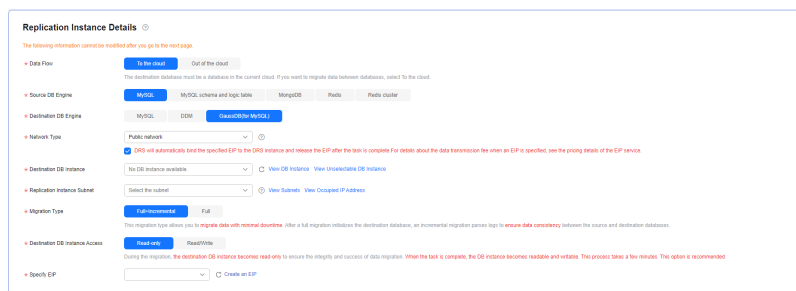


Table 2-32 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination is a DB instance on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select GaussDB(for MySQL) .
Network Type	Available options: VPC , Public network , and VPN or Direct Connect . By default, the value is Public network . <ul style="list-style-type: none"> – VPC is suitable for migrations between cloud databases of the same account in the same region. – VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. – Public network is suitable for migration from on-premises databases or external cloud databases to destination databases.
Destination DB Instance	Select the GaussDB(for MySQL) instance you have created.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"><li data-bbox="691 297 1428 465">– Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects in non-system databases are migrated to the destination database at one time, including tables, views, stored procedures, and triggers. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.<li data-bbox="691 633 1428 824">– Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.

Parameter	Description
Destination DB Instance Access	<ul style="list-style-type: none"> <li data-bbox="691 297 1426 752"> <p>– Read-only During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.</p> <p>If a migration task fails, the destination database restores to the read/write state. After the migration task is resumed, the destination database becomes read-only.</p> <p>The access settings of the destination instance are also affected by the access settings of the DB instance itself. Therefore, you are advised not to set the access settings of the DB instance on the RDS console.</p> <p>If the destination database is set to read-only for a migration task, you can disable read-only on the task details page. After read-only is disabled, the destination database cannot be set back to read-only. Even if the migration task fails and then resumed and started, the destination database cannot be set to read-only.</p> <li data-bbox="691 1010 1426 1451"> <p>– Read/Write During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks. Set the destination database to Read/Write only when you need to modify other data in the database during the migration.</p> <p>The task cannot be modified after being created.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 2-26 AZ




Table 2-33 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 2-27 Enterprise projects and tags




Table 2-34 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 projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database configuration

Figure 2-28 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 2-35 Source database settings

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

- Destination database configuration

Figure 2-29 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Migrate Definer to User Yes No

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

Table 2-36 Destination database settings

Parameter	Description
DB Instance Name	The GaussDB(for MySQL) instance selected during the migration 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
Migrate Definer to User	<p>Indicates whether to migrate the Definers of all source database objects to the destination database user entered during the connection test.</p> <ul style="list-style-type: none"> - Yes The Definers of all source database objects will be migrated to the user. Other users do not have permissions for database objects unless these users are authorized. For details about authorization, see How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration? <p>For example, if the view is <code>CREATE ALGORITHM=UNDEFINED DEFINER='username'@'%' SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> before migration,</p> <p>it is converted to <code>CREATE ALGORITHM=UNDEFINED DEFINER='drsUser'@'%' SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> after the migration.</p> <p>drsUser indicates the destination database user used for testing the connection.</p> <p>NOTE</p> <ul style="list-style-type: none"> ▪ For the migration from MySQL to GaussDB(for MySQL), you can only select Yes. After the migration, the Definers of all source database objects will be migrated to the user. ▪ For details about Definer, see the MySQL official document.

 **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 Task** page, select migration objects and click **Next**.

Figure 2-30 Migration object

Row Control Yes No

Migrate Incremental Accounts and Permissions Yes No
The migration of incremental accounts and permissions may fail because the source and destination MySQL versions and account encryption modes may be different.

Migrate Account Yes No
During a database migration, you need to separately migrate accounts and permissions. Certain accounts cannot be migrated to the destination database. Ensure that services are not affected.

<input checked="" type="checkbox"/> Account	Can Be Migrated	Permission	Password	Remarks
<input checked="" type="checkbox"/> root@% [%]	Yes	GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RE...	<input type="password"/>	-
<input checked="" type="checkbox"/> user@% [%]	Yes	GRANT SELECT, INSERT, UPDATE, DELETE ON **	<input type="password"/>	-
<input checked="" type="checkbox"/> root@% [%]	Yes	GRANT USAGE ON **	<input type="password"/>	-
<input checked="" type="checkbox"/> user_2_001@% [%]	Yes	GRANT SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, RE...	<input type="password"/>	-
<input checked="" type="checkbox"/> user_004_1@% [%]	Yes	GRANT USAGE ON **	<input type="password"/>	-
<input checked="" type="checkbox"/> user_005_003@% [%]	Yes	GRANT SELECT ON **	<input type="password"/>	-

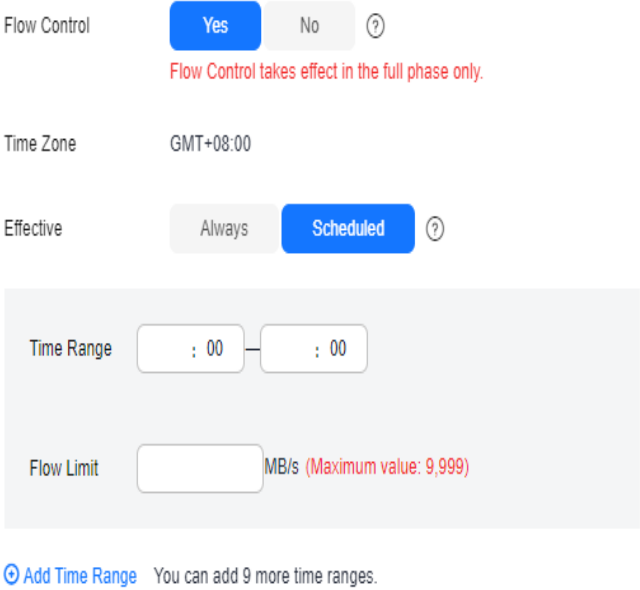
Real Password
 Set Unified Password


Create Indexes Along With Table Structure Yes No

Filter DROP DATABASE Yes No

Migration Object All Tables Databases

Table 2-37 Migrate 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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 2-31 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> - Flow control mode takes effect only during a full migration. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none"> • Yes If you need to migrate accounts, see Migrating Accounts. • No During migration, accounts, permissions, and passwords are not migrated.
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full migration phase.</p> <ul style="list-style-type: none"> • Yes: Indexes are migrated when the table structure is migrated in the full migration phase. • No: Indexes are migrated separately after data migration.
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> • All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. • Tables: The selected table-level objects will be migrated. • Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

NOTE

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

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

Figure 2-32 Task startup settings

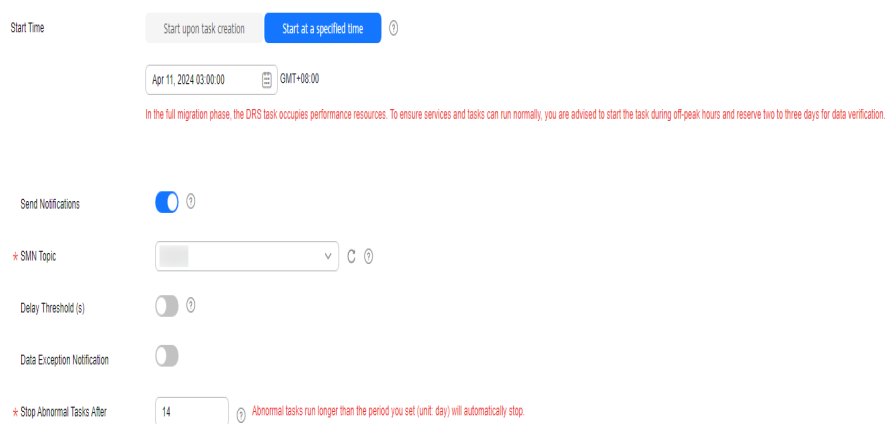



Table 2-38 Task startup settings

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

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.4 From MySQL Schema and Logic Table to DDM

Supported Source and Destination Databases

Table 2-39 Supported databases

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

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-40](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 2-40 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	The source sharded database user must have the following permissions: SELECT, SHOW VIEW, and EVENT	The source sharded database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT If the source database version is 8.0.2 or later, the XA_RECOVER_ADMIN permission is required to prevent data loss caused by uncommitted XA transactions during startup.
Destination database user	<ul style="list-style-type: none">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 migrated.	

 NOTE

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

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-41](#). DRS will automatically check the objects you selected before the migration.

Table 2-41 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none"> • The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: `!<>\/` • You can select table-level migration. • Only the source database data can be migrated. The table structure and other database objects of the source database cannot be migrated. • Only MyISAM and InnoDB tables can be migrated. • The system database and event statuses cannot be migrated. • Tables without primary keys cannot be migrated. • Non-standard floating-point data that can be written in loose mode but cannot be written in strict mode cannot be migrated. Such non-standard floating-point data may fail to be hit, causing data migration failures. • If the source database is MySQL 8.0, tables with function indexes cannot be migrated. • Full and incremental migrations do not support invisible columns. Invisible columns can be migrated 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> • If the source is a DDM instance, the tables cannot contain sharding keys of the timestamp type.

Precautions

To ensure tasks can run normally, DRS provides automatic pre-check. Before starting a DRS task, DRS checks the configurations and conditions of the source and destination databases. For details about the main check items and handling

suggestions, see [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 2-42](#).

Table 2-42 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The database names and table names of the source sharding middleware cannot contain non-ASCII characters or the following characters: .'<>\/• If the source database is an on-premises database and has Percona Server for MySQL 5.6.x or Percona Server for MySQL 5.7.x installed, the memory manager must use Jemalloc to prevent Out of Memory errors caused by frequent queries on system tables.• The source database does not support the mysql binlog dump command.• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.• 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.• Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database.

Type	Restrictions
Use restrictions	<p>General</p> <ul style="list-style-type: none"> • During the migration, do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. • During migration, do not clear the binlog in the source database. • During the migration, do not change the sharding key of a table on the source DDM instance, or change an unsharded or broadcast table to a sharded table, or change a sharded table to an unsharded or broadcast table. <p>Full migration</p> <ul style="list-style-type: none"> • During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> • During an incremental migration of table-level objects, renaming tables is not supported. • During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database. • During an incremental migration, 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. • During an incremental migration, the following DDL statements are supported: ADD COLUMN, CREATE TABLE, CREATE INDEX, DROP INDEX, RENAME INDEX and ADD INDEX. • During an incremental migration, perform DDL operations on the same table or column during off-peak hours at an interval of more than 1 minute. <p>Stopping a task</p> <ul style="list-style-type: none"> • 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.

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 migration task cannot be created.• If the data types are incompatible, the migration 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.• If a table in the source database contains system columns <code>_ddm_lock</code> and <code>_slot</code>, data in these columns will be filtered out during full migration and incremental migration, resulting in data loss.• After a task is created, you cannot add schemas to the source database or modify the old schema to associate with the new RDS DB instance. Otherwise, data cannot be migrated or the task fails.• If the destination DB instance uses columns of the <code>TIMESTAMP</code> or <code>DATETIME</code> data type as its sharding key, the seconds precision of the column is removed after the migration.

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 migration, 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](#).
- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

This section describes how to configure a task for migration from MySQL schema and logic table to DDM over a public network.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-33 Migration task information

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

Region: [Dropdown menu]

Project: [Dropdown menu]

* Task Name: DRS-5678

Description: [Text area]

0/256

Table 2-43 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-34 Replication instance information

Replication Instance Details

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

- Data Flow: **To the cloud** (Selected) / Out of the cloud
- Source DB Engine: MySQL / **MySQL, Oracle, and TSP Suite** (Selected) / MongoDB / Redis / Redis cluster
- Destination DB Engine: **DB2** (Selected)
- Network Type: Public network
- Destination DB Instance: **DB2 will automatically bind the specified EP to the DB2 instance and release the EP after the task is completed.** / No DB instance available / View DB Instance / View Unavailable DB Instance
- Application Instance Subnet: Select the subnet / View Subnets
- Migration Type: **Full Restoration** (Selected) / Full Incremental
- Source DB Instance Quantity: 2
- Specify EP: [Dropdown menu] / Create an EP

Table 2-44 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.

Parameter	Description
Source DB Engine	Select MySQL schema and logic table .
Destination DB Engine	Select DDM .
Network Type	<p>The Public network is used as an example. Available options: VPC, VPN or Direct Connect, and Public network. By default, the value is Public network.</p> <ul style="list-style-type: none"> - VPC is suitable for migrations between cloud databases of the same account in the same region and VPC. - Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases. - VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you have created.
Replication Instance Subnet	<p>The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects in non-system databases are migrated to the destination database at one time, including tables, views, stored procedures, and triggers. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database. - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.
Source DB Instance Quantity	Specifies the number of DB instances bound to the source DDM database. The default value is 2 . The value ranges from 1 to 64 . Set this parameter based on the site requirements.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- AZ

Figure 2-35 AZ




Table 2-45 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 2-36 Enterprise projects and tags

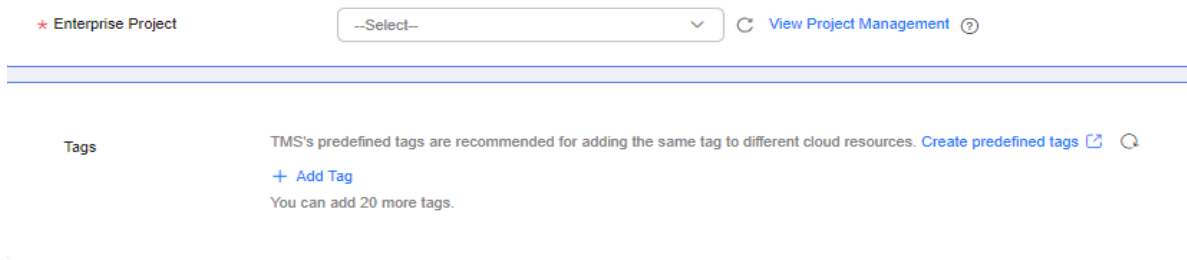


Table 2-46 Enterprise Project and Tags

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 projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 2-37 Source database information

Source Database

Middleware IP Address or Domain Name

Port

Middleware Username

Database Password

DB Instance

IP Address or Domain Name	Port	Username	Password
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="password"/>

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

Table 2-47 Source database information

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 for accessing the source database.
Middleware 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.
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.

- Destination database configuration

Figure 2-38 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Table 2-48 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration 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.

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 Task** page, select migration objects and click **Next**.

Figure 2-39 Migration object

Flow Control

Migration Object

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

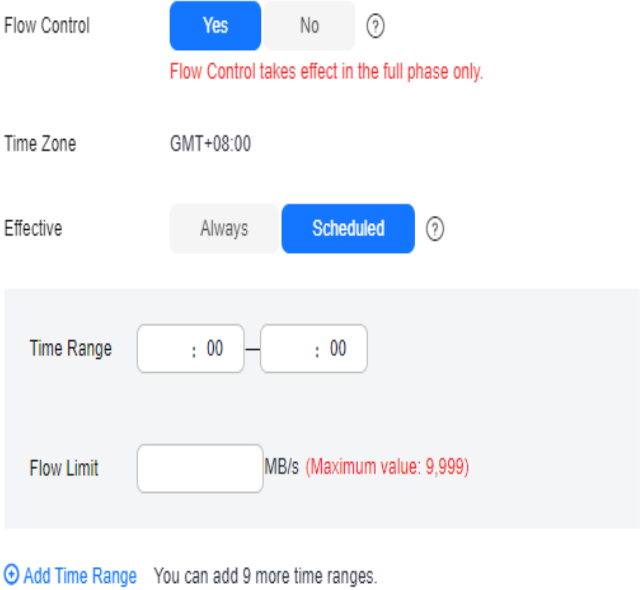
Search the expanded database using regular expressions:

Select All

Search the expanded database using regular expressions:

db_tatus_check_009 database

Table 2-49 Migrate 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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 2-40 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> Flow control mode takes effect only during a full migration. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. After the selected objects are migrated to the destination database, the object names will remain the same as those in the source database and cannot be modified.</p> <p>NOTE</p> <ul style="list-style-type: none"> Structures that are not created in the destination database are not migrated by default. 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 migration object cannot contain spaces. To quickly select the desired database objects, you can use the search function.

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

- If any check fails, 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*.

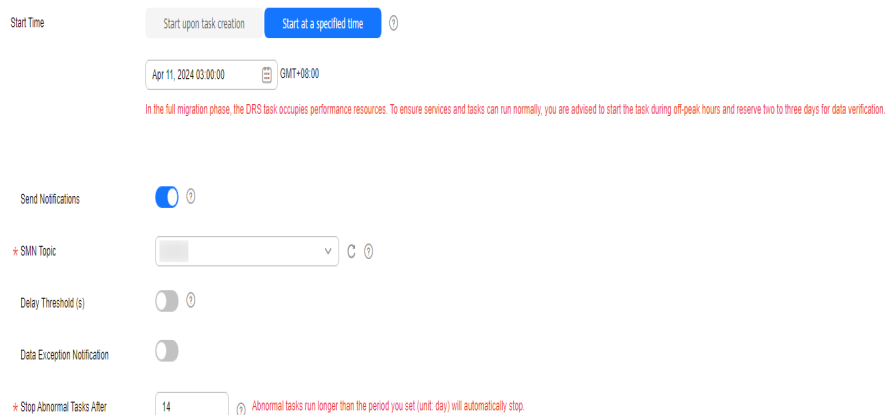
- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

You can proceed to the next step only when all checks are successful. If any alarms are generated, view and confirm the alarm 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** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 2-41 Task startup settings



Start Time

Start upon task creation Start at a specified time ⓘ

Apr 11, 2024 03:00:00 GMT+08:00

In the full migration phase, the DRS task occupies performance resources. To ensure services and tasks can run normally, you are advised to start the task during off-peak hours and reserve two to three days for data verification.

Send Notifications ⓘ

* SMN Topic ⓘ

Delay Threshold (s) ⓘ


Data Exception Notification

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

Table 2-50 Task startup settings

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.5 From MongoDB to DDS

Supported Source and Destination Databases

Table 2-51 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0)• ECS-hosted MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0)• Other Cloud MongoDB 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0• DDS DB instances (versions 3.2, 3.4, 4.0, 4.2, 4.4 and 5.0) <p>NOTE DDS 5.0 supports replica sets only. If the source database is a DDS DB instance, the source DB engine is DDS. Otherwise, the source DB engine is MongoDB-DDS.</p>	<ul style="list-style-type: none">• DDS DB instances (versions 3.4, 4.0, 4.2, 4.4, and 5.0) <p>NOTE The destination database version must be the same as or later than the source database version. DDS 5.0 supports replica sets only.</p>

You can run the following command to query a database version.

```
db.version()
```

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-52](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 2-52 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	<ul style="list-style-type: none"> ● Replica set: The source database user must have the readAnyDatabase permission for the admin database. ● Single node: The source database user must have the readAnyDatabase permission for the admin database. ● Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. ● To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database. 	<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 mongos node user must have the readAnyDatabase permission for the admin database and the read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database and the read permission for the local database. ● To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database.
Destination database user	<p>The user who connects to the destination database must have the dbAdminAnyDatabase permission of the admin database and the readWrite permission of the destination database.</p> <p>If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.</p>	

 **NOTE**

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

```
db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"},  
{role:"read",db:"config"}])
```

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-53](#). DRS will automatically check the objects you selected before the migration.

Table 2-53 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none"> • You can select table-, database-, or instance-level (all-level) migration. • Replica set: Only collections (including validator and capped collections), indexes, and views can be migrated. • Cluster: Only collections (including validator and capped collections), shard keys, indexes, and views can be migrated. • Single node: Only collections (including validator and capped collections), indexes, and views can be migrated. • Only user data and source database account information can be migrated. The system databases (for example, local, admin, and config) and system collection cannot be migrated. If service data is stored in the system database, run the renameCollection command 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 ranged sharding is used, maxKey cannot be used as the primary key. • 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 migration. • If the source database is a cluster instance of version 4.4 or later and the destination database version is earlier than 5.0, composite hash shard keys and composite hash indexes are not supported. If the destination database version is 5.0, composite hash shard keys and composite hash indexes are supported. • If the source database is a replica set instance of version 4.4 or later and the destination database version is earlier than 5.0, composite hash indexes are not supported. If the destination database version is 5.0, composite hash indexes are supported. • Time series collections are not supported in both the full and incremental phases.

Network Connection Between the Source and Destination Databases

- If the source database is deployed on a local host, select a network type by referring to [From On-premises Databases to Huawei Cloud](#).
- If the source database is a database of other cloud, select a network type by referring to [From Other Cloud Databases to Huawei Cloud](#).

- If the source database is an ECS-hosted database, select a network type by referring to [From ECS-Hosted Databases on Huawei Cloud to Huawei Cloud](#).
- If the source database is a Huawei Cloud database, select a network type by referring to [From Huawei Cloud to Huawei Cloud](#).

Log Configuration Check on the Source Database

During an incremental migration, Olog of the source database must be enabled. If the storage space is sufficient, store the source database Olog for as long as possible. The recommended retention period is at least three days.

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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 2-54](#).

Table 2-54 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source cannot be a GeminiDB Mongo instance.• The source database name cannot contain /\.\$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).• When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique.• Replica set: The MongoDB replica set instance must be available and have primary nodes.• If you need to perform an incremental migration for a single node instance, the source database must be a DDS single-node instance on the current cloud.

Type	Restrictions
Restrictions on usage	<p>General</p> <ul style="list-style-type: none"> • During the migration, do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. • During migration, data rollback caused by a primary/standby switchover of the source database is not supported. • During the migration, documents larger than 16 MB in the source database cannot be inserted or updated. • During the migration, do not run sh.moveChunk() on the source database. Otherwise, the migrated data will be inconsistent. <p>Full migration</p> <ul style="list-style-type: none"> • During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> • During an incremental migration of collections, you are advised not to rename the collections. • To ensure the migration performance, concurrent replay is performed at the set level in the incremental task phase. In the following scenarios, only single-thread write is supported and concurrent replay is not supported: <ul style="list-style-type: none"> - The collection index contains a unique key. - The value of capped of the collection attribute is true. <p>In either of the preceding scenarios, the task delay may increase.</p> <p>Data comparison</p> <ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • During row comparison, if an orphan document exists in a cluster instance or chunks are being migrated, the number of returned rows is incorrect and the comparison results are inconsistent. For details, see MongoDB official documents. <p>Troubleshooting</p>

Type	Restrictions
	<ul style="list-style-type: none">• If any problem occurs during task creation, startup, full synchronization, incremental synchronization, or completion, rectify the fault by referring to Troubleshooting.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If the source database is not on a cluster instance, the following operations and commands are supported during incremental migration: <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. ● During a full+incremental migration between clusters or from a cluster to a replica set, the objects to be migrated cannot be deleted. Otherwise, the migration task fails. ● If you select Cluster (MongoDB 4.0+) for Source DB Instance Type, DRS will use the MongoDB change streams feature during the migration. Note the following before you use change streams: <ul style="list-style-type: none"> - Data subscription using change streams consumes a certain amount of CPU and memory resources of the source database. Evaluate the resources of the source database 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. - Change streams support only the following DDLs: drop database, drop collection and rename - The <code>DBPointer</code> and <code>DBRef</code> data types are not supported. - In the incremental migration phase, the migration speed can reach up to 10,000 rows in a single table per second. - Currently, only whitelisted users can use Change Streams. ● If you select <code>Olog</code> for <code>Source Database Type</code>, DRS will create multiple subtasks based on the number of source shards. If <code>Task Rate Limit</code> is selected, the configured rate limit value is synchronized to each subtask. ● If a <code>Time-to-Live (TTL)</code> index already exists in the collection of the source database or is created during an incremental migration, data consistency cannot be ensured when source and destination databases are in different time zones or clocks. ● 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

Type	Restrictions
	<p>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 migrate compression parameters because the destination database does not support compression parameter settings.</p> <ul style="list-style-type: none"> ● If the accounts and roles to be migrated conflict with those in the destination database, DRS will skip the conflict data and continue the migration. ● 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. ● If the source is a replica set instance, enter information about all primary and secondary nodes to reduce the impact of a primary/secondary switchover on the migration task. If you enter information about multiple primary and secondary nodes, ensure that all nodes belong to the same replica set instance. ● If the source is a cluster instance, enter information about multiple mongos nodes to reduce the impact of single-node failure on the migration task. In addition, ensure that all mongos nodes belong to the same cluster instance. For an incremental migration of a cluster instance, you are advised to enter information about all primary and secondary nodes of the shard node and ensure that all node information belongs to the same shard to reduce the impact of a primary/secondary switchover on the migration task. Ensure that all shard nodes belong to the same cluster. ● In some migration scenarios, to prevent the drop database operation from deleting the existing collections in the destination database, the drop database operation will not be synchronized to the destination database. <ul style="list-style-type: none"> - If the source database version is earlier than MongoDB 3.6, running the drop database command will delete the collections only from the source database. The collections in destination database will not be deleted. - If the source database version is MongoDB 3.6 or later, the drop database operation is represented by the drop database and drop collection operations in oplog. Running the drop database command will delete the collections from both the source and destination databases. ● To accelerate the migration, delete unnecessary indexes from the source database and retain only necessary indexes before the migration. You are advised not to create indexes for the source database during the migration. If indexes must be created, create them in the background.

Type	Restrictions
	<ul style="list-style-type: none"> To prevent loopback, do not start tasks that migrate the same database to and out of the cloud at the same time.

Procedure

This section uses a MongoDB sharded cluster as an example to describe how to configure a task for migrating data from MongoDB to DDS Community Edition over a public network.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-42 Migration task information

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

Region

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

Project

* Task Name

Description

0/256

Table 2-55 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-43 Replication instance information

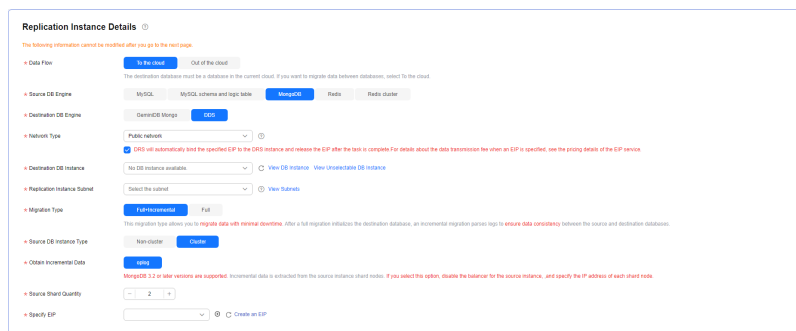


Table 2-56 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.
Source DB Engine	Select MongoDB .
Destination DB Engine	Select DDS .
Network Type	Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network . <ul style="list-style-type: none"> – VPC is suitable for migrations between cloud databases of the same account in the same region and VPC. – Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases. – VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you have created.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is permitted. It migrates all objects and data in non-system databases to the destination database at one time. The objects include collections, views, and indexes. <p>NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.</p> <ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
Source DB Instance Type	<p>If you select Full+Incremental for Migration Type, set this parameter based on the source database.</p> <ul style="list-style-type: none"> - If the source database is a cluster instance, set this parameter to Cluster. - If the source database is a replica set or a single node instance, set this parameter to Non-cluster.

Parameter	Description
Obtain Incremental Data	<p>If Source DB Instance Type is set to Cluster and Migration Mode is set to Full+Incremental, you need to determine how to capture data changes during the incremental synchronization. Also, the balancer of the source database must be disabled and orphaned documents must be deleted. For details, see How Do I Disable the Balancer? and How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?</p> <ul style="list-style-type: none"> – oplog: For MongoDB 3.2 or later, DRS directly connects to each shard of the source DB instance to extract data. If you select this method, you need to enter the connection information of each shard in the source database when testing the connection between the source database and the DRS instance. – changeStream: This method is recommended. For MongoDB 4.0 and later, DRS connects to mongos nodes of the source database to extract data. If you select this method, you must enable the WiredTiger storage engine of the source database. <p>NOTE Only whitelisted users can use changeStream. 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.</p>
Source Shard Quantity	<p>If Source DB Instance Type is set to Cluster and Obtain Incremental Data is set to oplog, you need to enter the number of source database shards.</p> <p>The number of source shards ranges from 2 to 64. Specify this parameter based on the actual number of shards in the source DB.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. If Source DB Instance Type is set to Cluster and Obtain Incremental Data is set to oplog, the number of specified EIPs must be consistent with that of source shards.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 2-44 AZ




Table 2-57 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 2-45 Enterprise projects and tags

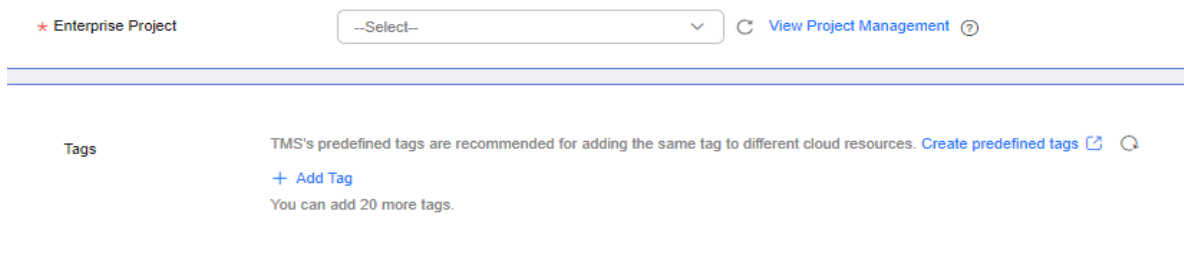


Table 2-58 Enterprise Project and Tags

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 projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 2-46 Source database information

Table 2-59 Source database information

Parameter	Description
mongos IP Address or Domain Name	<p>IP address or domain name of the source database in the IP address/Domain name:Port format. The port of the source database. Range: 1 - 65535</p> <p>If the source database is an on-premises MongoDB database, enter the IP address or domain name of the mongos node.</p> <p>You can enter a maximum of three groups of IP addresses or domain names of the source 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 sharded cluster.</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 authentication database of the source MongoDB database user. For example: The default authentication database of DDS instance is admin .
mongos Username	The username of the mongos node of the source MongoDB database. For example, the default mongos username of a DDS instance on Huawei Cloud is rwuser .
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.
Sharded Database	<p>Enter shard node information of the MongoDB database based on the number of cluster shard nodes in the source database.</p> <ul style="list-style-type: none">- Set the IP address or domain name of the sharded database to that of the shard node in the source MongoDB database. The value is in the IP address/Domain name:Port format. If the source database is a Huawei Cloud DDS instance, enable IP addresses of shard nodes by referring to Enabling IP Addresses of Shard and Config Nodes and enter the shard IP information.- Set the authentication database of the sharded database to the name of the database to which the shard node account of the source MongoDB database belongs. For example, the default shard authentication database of a DDS instance on Huawei Cloud is admin.- Set the username of the sharded database to that of the shard node in the source MongoDB database. For example, the default shard username of a DDS instance on Huawei Cloud is sharduser.

 **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.

- Destination database configuration

Figure 2-47 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 2-60 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration task. This parameter cannot be changed.
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is admin .
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is disabled, your data may be at risk.

NOTE

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

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

Figure 2-48 Migration object


Flow Control: Yes **No** ⓘ

Migrate Account: **Yes** No
Confirm All Remarks

Account Information ⓘ

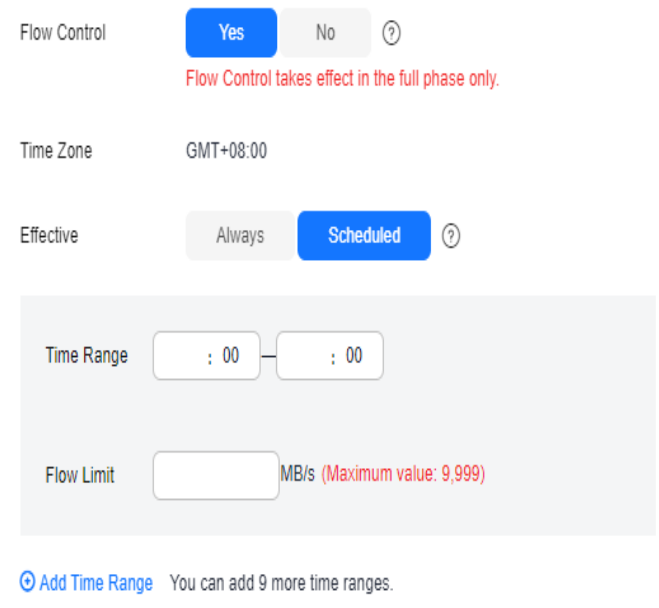
Account	Can Be Migrated	Role	Remarks
<input type="checkbox"/> admin.rwuser	No	admin.root	View Confirm


Role Information

Role Name	Can Be Migrated	Permission	Inherited Role	Remarks
 No data available.				

Migration Object: **All** Tables Databases

Table 2-61 Migrate 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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 2-49 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> - Flow control mode takes effect only during a full migration. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Account	<p>Accounts to be migrated can be classified into the following types: accounts that can be migrated and accounts that cannot be migrated. You can choose whether to migrate the accounts. Accounts that cannot be migrated or accounts that are not selected will not exist in the destination database. Ensure that your services will not be affected by these accounts.</p> <ul style="list-style-type: none"> • Yes If you need to migrate accounts, see Migrating Accounts. • No During the migration, accounts and roles are not migrated.
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> • All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. • Tables: The selected table-level objects will be migrated. • Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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


Figure 2-50 Task startup settings

Table 2-62 Task startup settings

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

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.6 From MongoDB to GeminiDB Mongo

Supported Source and Destination Databases

Table 2-63 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises MongoDB (versions 3.2, 3.4, and 4.0)Self-built MongoDB on ECS (versions 3.2, 3.4, and 4.0)MongoDB 3.2, 3.4, and 4.0 on other clouds	<ul style="list-style-type: none">GeminiDB Mongo instances (versions 3.4 and 4.0) <p>NOTE The destination database version must be the same as or later than the source database version.</p>

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-64](#). DRS will automatically check the objects you selected before the migration.

Table 2-64 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">You can select table-, database-, or instance-level (all-level) migration.Replica set: Only collections (including validator and capped collections), indexes, and views can be migrated.Cluster: Only collections (including validator and capped collections), shard keys, indexes, and views can be migrated.Single node: Only collections (including validator and capped collections), indexes, and views can be migrated.The statement for creating a view cannot contain a regular expression.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.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 migration.

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-65](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 2-65 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	<ul style="list-style-type: none"> • Replica set: The source database user must have the readAnyDatabase permission for the admin database. • Single node: The source database user must have the readAnyDatabase permission for the admin database. • Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. • To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database. 	<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 mongos node user must have the readAnyDatabase permission for the admin database and the read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database and the read permission for the local database. • To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database.

Type	Full Migration	Full+Incremental Migration
Destination database user	<p>The user who connects to the destination database must have the dbAdminAnyDatabase permission of the admin database and the readWrite permission of the destination database.</p> <p>If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.</p>	

 **NOTE**

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

```
db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"},
{role:"read",db:"config"}])
```

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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 2-66](#).

Table 2-66 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none"> • The source cannot be a GeminiDB Mongo instance. • The source database name cannot contain /\. \$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$). • When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique. • Replica set: The MongoDB replica set instance must be available and have primary nodes.

Type	Restrictions
Use restrictions	<p>General</p> <ul style="list-style-type: none"> • During the migration, do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. • During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. • During migration, data rollback caused by a primary/standby switchover of the source database is not supported. • During the migration, documents larger than 16 MB in the source database cannot be inserted or updated. • During the migration, do not run sh.moveChunk() on the source database. Otherwise, the migrated data will be inconsistent. <p>Full migration</p> <ul style="list-style-type: none"> • During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> • During an incremental migration of collections, you are advised not to rename the collections. <p>Data comparison</p> <ul style="list-style-type: none"> • You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. • During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. • During row comparison, if an orphan document exists in a cluster instance or chunks are being migrated, the number of returned rows is incorrect and the comparison results are inconsistent. For details, see MongoDB official documents.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none">• If the source database is not on a cluster instance, the following operations and commands are supported during incremental migration:<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.• During a full plus incremental migration between clusters, the objects to be migrated cannot be deleted. Otherwise, the migration task fails.• If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental migration, data consistency cannot be ensured when source and destination databases are in different time zone.• If the destination database does not support compression parameter settings, DRS does not migrate compression parameters, and no error is reported during the migration.• If the MongoDB service of the source database is deployed with other services on the same server, set the value of the <code>cacheSizeGB</code> parameter to the half of the minimum idle cache for the WiredTiger engine of the source database.• If the source is a replica set instance, enter information about all primary and secondary nodes to reduce the impact of a primary/secondary switchover on the migration task. If you enter information about multiple primary and secondary nodes, ensure that all nodes belong to the same replica set instance.

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 migration, 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](#).
- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

This section describes how to use DRS to migrate data from MongoDB to GeminiDB Mongo over a public network.

Step 1 On the **Create Replication Instance** page, select a region and project, specify the task name, description, and the replication instance details, and click **Create Now**.

- Task information description

Figure 2-51 Migration task information

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

Region

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

Project

* Task Name

Description

0/256

Table 2-67 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-52 Replication instance information

Replication Instance Details

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

• Data Flow In-Cloud Out of the cloud

The destination database must be a database in the current cloud. If you want to migrate data between databases, select To the cloud.

• Source DB Engine MySQL MySQL schemas and logs table MongoDB Redis Redis cluster

• Destination DB Engine GaussDB MySQL DDS

• Network Type Cloud network VPC

DDS will automatically bind the specified EIP to the DDS instance and release the EIP after the task is completed. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance No DB instance available View DB Instance View Unconnected DB Instance

• Replication Instance Subnet Select the subnet View Subnets

• Migration Type Full-Incremental Full

This migration type allows you to migrate data with minimal downtime. After a full migration releases the destination database, an incremental migration parses logs to ensure data consistency between the source and destination databases.

• Source DB Instance Type Non-Cluster Cluster

• Source Shard Quantity

• Standby EIP

Table 2-68 Replication 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 GeminiDB Mongo .
Network Type	Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network . <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region and VPC.- Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases.- VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you have created.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is permitted. It migrates all objects and data in non-system databases to the destination database at one time. The objects include collections, views, and indexes. <p>NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.</p> <ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
Source DB Instance Type	<p>If you select Full+Incremental for Migration Type, set this parameter based on the source database.</p> <ul style="list-style-type: none"> - If the source database is a cluster instance, set this parameter to Cluster. - If the source database is a replica set or a single node instance, set this parameter to Non-cluster.
Source Shard Quantity	<p>If Source DB Instance Type is set to Cluster, you need to enter the number of source database shards.</p> <p>The number of source shards ranges from 2 to 32. Specify this parameter based on the actual number of shards in the source DB.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 2-53 AZ

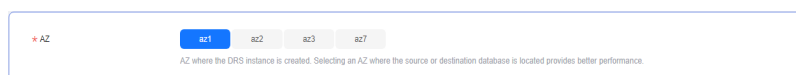


Table 2-69 Task AZ

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 2-54 Enterprise projects and tags

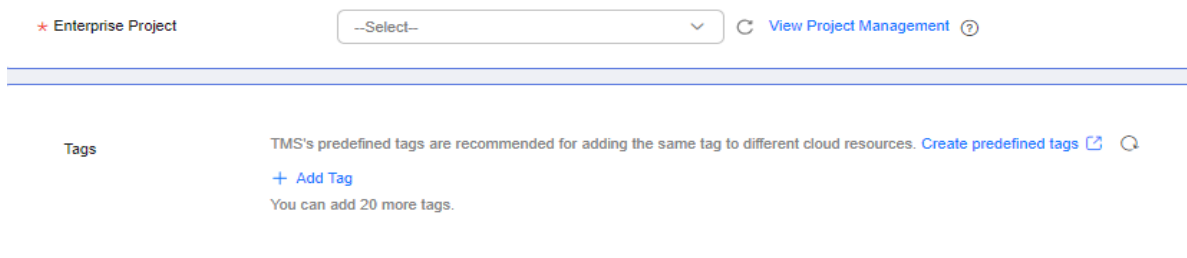


Table 2-70 Enterprise Project and Tags

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 projects, 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 2 On the **Configure Source and Destination Databases** page, wait until the replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Figure 2-55 Source database information

Source Database

IP Address or Domain Name ⓘ
Ensure that the entered addresses belong to the same DB instance.

Authentication Database

Database Username

Database Password ⓘ

SSL Connection

Table 2-71 Source database information

Parameter	Description
IP Address or Domain Name	<p>IP address or domain name of the source database in the IP address/Domain name:Port format. The port of the source database. Range: 1 - 65535</p> <p>You can enter a maximum of three groups of IP addresses or domain names of the source 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 sharded cluster.</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 authentication database.
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.

- Destination database configuration

Figure 2-56 Destination database information

Destination Database

DB Instance Name

Authentication Database

Database Username

Database Password

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

Table 2-72 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration task. This parameter cannot be changed.
Authentication Database	The name of the authentication database.
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 replication instance during the migration. After the task is deleted, the username and password are permanently deleted.

Step 3 On the **Set Task** page, select migration objects and click **Next**.

Figure 2-57 Migration object

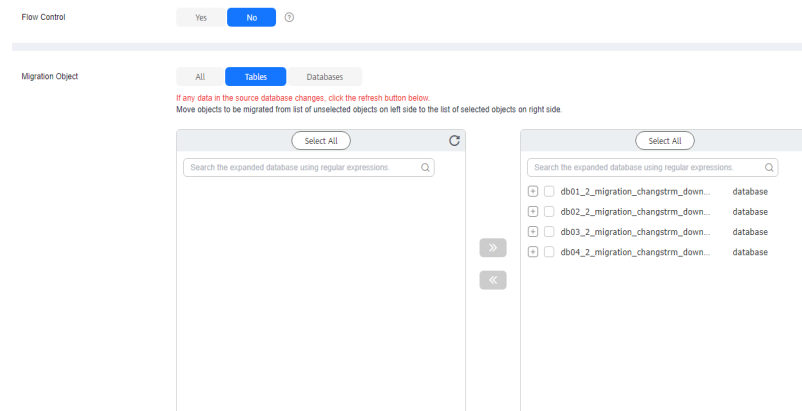
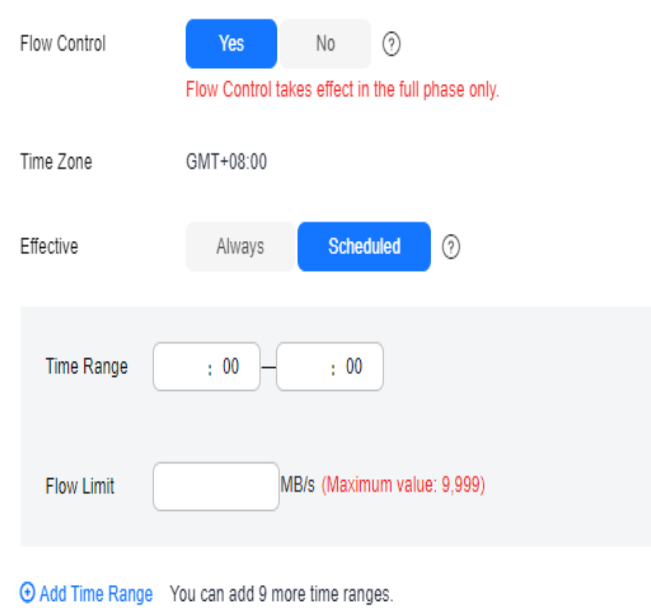



Table 2-73 Migrate Object

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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 2-58 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> Flow control mode takes effect only during a full migration. You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> • All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. • Tables: The selected table-level objects will be migrated. • Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as views, in the database to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

Step 4 On the **Check Task** page, check the migration task.

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Step 5 On the displayed page, specify **Start Time**, **Send Notifications**, **SMN Topic**, **Delay Threshold (s)**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 2-59 Task startup settings

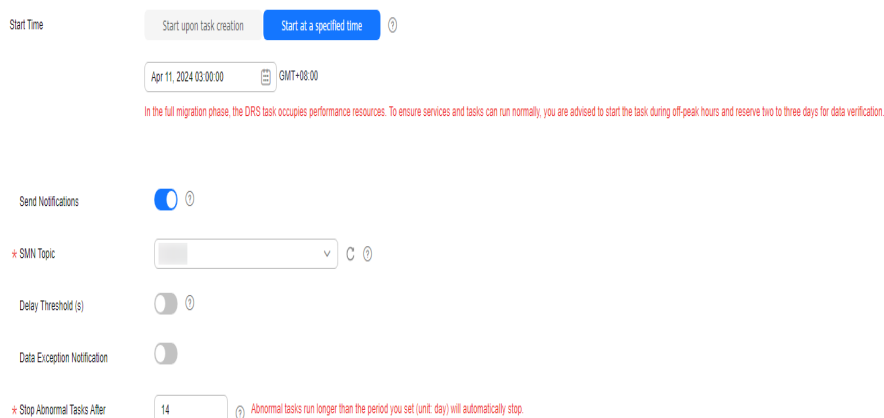



Table 2-74 Task startup settings

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

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

Step 6 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.7 From Redis to GeminiDB Redis

Supported Source and Destination Databases

Table 2-75 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises official open-source single-node or primary/standby Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x)ECS-hosted official open-source single-node or primary/standby Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x)	GeminiDB Redis NOTE The destination database version must be the same as or later than the source database version.

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-76](#). DRS will automatically check the objects you selected before the migration.

Table 2-76 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">Object level: database levelSupported migration objects:<ul style="list-style-type: none">Data of the String, Hash, List, Set, and Sorted Set types can be migrated.Lua scripts can be migrated (except in a single full migration task).Transactions cannot be migrated.Data of the Stream, BitMap, HyperLogLog, and GEO types cannot be migrated.User-defined types are not supported.Commands that are not supported by the destination GeminiDB Redis instance cannot be migrated.

Suggestions

⚠ CAUTION

To maintain data consistency before and after the migration, ensure that no data is written to your source and destination databases during a full migration. In the full+incremental migration mode, you can continue the migration while data is still being written to the source database.

- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
- In the migration, ensure that no data is written to the destination database to ensure data consistency before and after the migration.
- Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full migration, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental migration process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full migration task contains three phases. To ensure smooth migration, read the following precautions before creating a migration task.

Table 2-77 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - To create a full+incremental task, the PSYNC command must be supported. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination must be a GeminiDB Redis instance on the current cloud. - The destination GeminiDB Redis instance must be empty before the migration. - The version of the destination database instance must be the same as or later than that of the source database. ● Other notes: <ul style="list-style-type: none"> - A full migration task reads data in SCAN mode. - A full+incremental task read data by parsing RDB files. - If the source database does not support the PSYNC command, use the full migration type. - If the source database supports the PSYNC command, use the full+incremental migration type. - The expiration time settings of the full migration task cannot be migrated. - In a full migration task, Lua scripts cannot be migrated. - When migrating Lua scripts, ensure that the source database contains other data. Otherwise, Lua scripts in RDB cannot be parsed, causing the Lua script migration failure. - If the source is a self-managed single-node Redis database, the following data types can be migrated during the incremental migration phase: append, blpop, brpop, decr, decrby, del, expire, expireat, flushall, flushdb, getset, hdel, hincrby, hincrbyfloat, hmset, hset, hsetnx, incr, incrby, incrbyfloat, linsert, lpop, lpush, lpushx, lrem, lset, ltrim, move, mset, msetnx, multi, persist, pexpire, pexpireat, psetex, publish, rename, rpop, rpush, rpushx, sadd, select, set, setbit, setex, setnx, setrange, smove, spop, srem, zadd, zincrby, zrem, zremrangebylex, zremrangebyrank and zremrangebyscore. - GeminiDB Redis does not support multi-key commands. If the source is a standalone Redis database and keys involved in an operation aren't in the same hash-slot, the following error message will appear: CROSSSLOT Keys in request don't hash to the same slot - DRS does not perform high-risk operations such as FLUSH on the destination database, so data of the List type may be duplicated.

Type	Restrictions
	<ul style="list-style-type: none">- To ensure a successful incremental migration, increase the value of repl-backlog-size in the redis.conf configuration file of the source database before the migration.- During the migration, if the self-built Redis database is scaled in or out (for example, shards are added or deleted) or the specifications are changed (for example, the memory is scaled up), reconfigure the migration task. To ensure data consistency, clear the data that has been migrated to the destination GeminiDB Redis before reconfiguring the task.- If the expire policy is configured for some keys in the source database, the number of keys in the destination database (for example, by running the info command) may be less than that in the source database because the keys may have expired but are not deleted. To ensure data consistency, you are advised to prolong the key expiration time.
Full migration	<ul style="list-style-type: none">• During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases.• During the migration, do not modify the destination database that is being migrated.
Incremental migration	<ul style="list-style-type: none">• During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases.• During the migration, do not modify the destination database that is being migrated.

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 migration, 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](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section describes how to use DRS to migrate data from a Redis database to GeminiDB Redis over a public network.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-60 Migration task information

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

Region

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

Project

* Task Name

Description

0/256

Table 2-78 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-61 Replication instance information

Replication Instance Details

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

Data Flow: To the cloud Out of the cloud

The destination database must be a database in the current cloud. If you want to migrate data between databases, select To the cloud.

Source DB Engine: MySQL, schema and topic table MergeDB

Destination DB Engine:

Network Type:

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

Destination DB Instance:

Replication Instance Subnet:

Migration Type: Full Full

This migration type allows you to migrate data with minimal downtime. After a full migration releases the destination database, an incremental migration process logs to ensure data consistency between the source and destination databases.

Specify EP:

Table 2-79 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.

Parameter	Description
Source DB Engine	Select Redis .
Destination DB Engine	Select GeminiDB Redis .
Network Type	<p>Available options: VPC, VPN or Direct Connect, and Public network. Public network is selected in this example.</p> <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region and VPC.- Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases.- VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you created.
Replication Instance Subnet	<p>The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is permitted. It migrates all objects and data in non-system databases to the destination database at one time. The objects include collections, views, and indexes. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database. - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 2-62 AZ




Table 2-80 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 2-63 Enterprise projects and tags

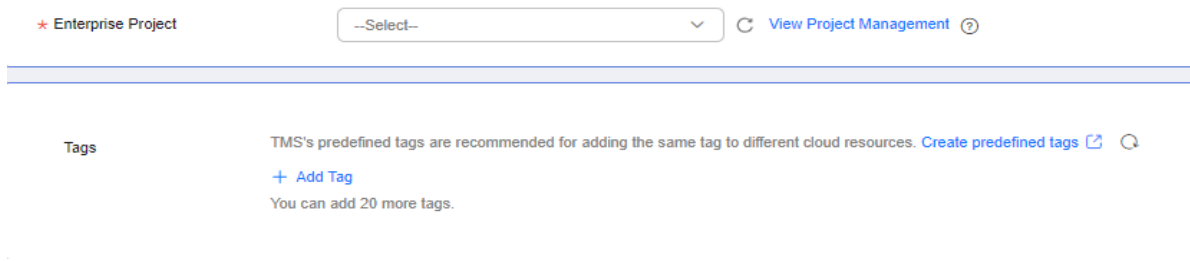


Table 2-81 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise projects, 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 replication instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 2-64 DNS Server

Configure Your Own DNS Server ⓘ

DNS Server

DNS Server IP Address . .

Table 2-82 DNS server information

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

NOTE

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

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

- Source database information

Figure 2-65 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.

IP Address or Domain Name

Port

Database Password

Table 2-83 Source database settings

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

Parameter	Description
Database Password	Password used to log in to the source database.

 **NOTE**

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

- Destination database configuration

Figure 2-66 Destination database information

Destination Database

DB Instance Name

Database Password

Table 2-84 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration task. This parameter cannot be changed.
Database Password	The password for the database username.

 **NOTE**

The password of the destination database is encrypted and stored in the database and the replication instance during the migration. After the task is deleted, the password is permanently deleted.

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

Figure 2-67 Migration object

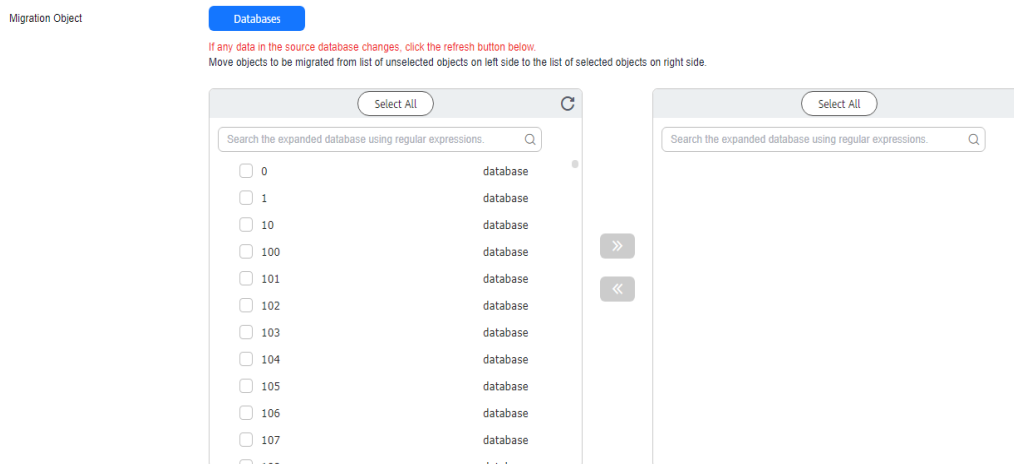



Table 2-85 Migration object

Parameter	Description
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. DRS supports database-level migration. You can select the databases to be migrated based on service requirements.</p> <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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

Figure 2-68 Task startup settings

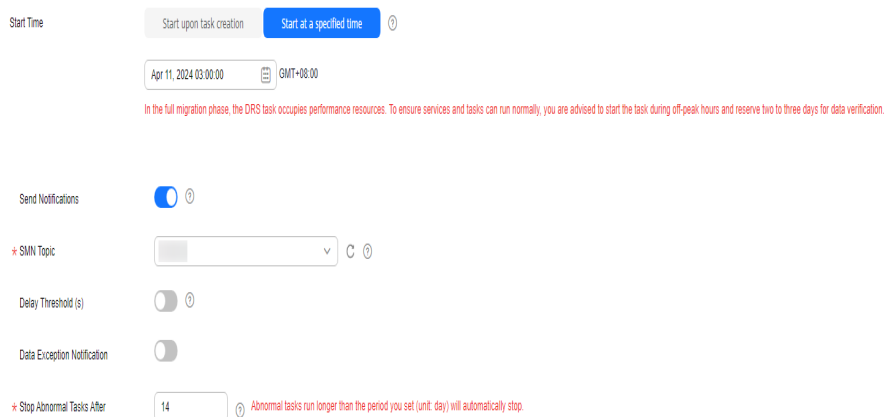



Table 2-86 Task startup settings

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

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

2.8 From Redis Cluster to GeminiDB Redis

Supported Source and Destination Databases

Table 2-87 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises open-source Redis Cluster (versions 3.0.x, 4.0.x, and 5.0.x)ECS-hosted open-source Redis Cluster (versions 3.0.x, 4.0.x, and 5.0.x)	GeminiDB Redis NOTE The destination database version must be the same as or later than the source database version.

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-88](#). DRS will automatically check the objects you selected before the migration.

Table 2-88 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">Object level: all database objectsSupported migration objects:<ul style="list-style-type: none">Data of the String, Hash, List, Set, and Sorted Set types can be migrated.Lua scripts can be migrated.Transactions cannot be migrated.Data of the Stream, BitMap, HyperLogLog, and GEO types cannot be migrated.User-defined types are not supported.Commands that are not supported by the destination GeminiDB Redis instance cannot be migrated.

Suggestions

⚠ CAUTION

To maintain data consistency before and after the migration, ensure that no data is written to your source and destination databases during a full migration. In the full+incremental migration mode, you can continue the migration while data is still being written to the source database.

- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
- In the migration, ensure that no data is written to the destination database to ensure data consistency before and after the migration.
- Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full migration, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental migration process consists of four phases: task startup, full synchronization, incremental synchronization, and task completion. A single full migration task contains three phases. To ensure smooth migration, read the following precautions before creating a migration task.

Table 2-89 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database object requirements: <ul style="list-style-type: none"> - To create a full+incremental task, the PSYNC command must be supported. ● Destination database object requirements: <ul style="list-style-type: none"> - The destination must be a GeminiDB Redis instance on the current cloud. - It is recommended that the destination GeminiDB Redis instance be empty before the migration. - The version of the destination database instance must be the same as or later than that of the source database. ● Other notes: <ul style="list-style-type: none"> - A full+incremental task read data by parsing RDB files. - When migrating Lua scripts, ensure that the source database contains other data. Otherwise, Lua scripts in RDB cannot be parsed, causing the Lua script migration failure. - DRS does not perform high-risk operations such as FLUSH on the destination database, so data of the List type may be duplicated. - When testing the connection to the source database, you must enter the IP addresses of all shards in the Redis Cluster. To keep services accessible to the source Redis instance, you are advised to migrate data from the replica node (slave node) of the source instance. That is, you are advised to enter the IP address of the slave node in the Redis Cluster instance. - To ensure a successful incremental migration, increase the value of repl-backlog-size in the redis.conf configuration file of the source database before the migration. - During the migration, if the source Redis Cluster instance is scaled in or out (for example, shards are added or deleted) or the specifications are changed (for example, the memory is scaled up), reconfigure the migration task. To ensure data consistency, clear the data that has been migrated to the destination GeminiDB Redis before reconfiguring the task. - If the expire policy is configured for some keys in the source database, the number of keys in the destination database (for example, by running the info command) may be less than that in the source database because the keys may have expired but are not deleted. To ensure data consistency, you are advised to prolong the key expiration time.

Type	Restrictions
Full migration	<ul style="list-style-type: none">• During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases.• During migration, if a primary/standby switchover occurs in the source Redis Cluster instance, the task will fail. In this case, you need to clear the destination database to reconfigure the task.• During the migration, do not modify the destination database that is being migrated.
Incremental migration	<ul style="list-style-type: none">• During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases.• During migration, if a primary/standby switchover occurs in the source Redis Cluster instance, the task will fail. In this case, you need to clear the destination database to reconfigure the task.• During the migration, do not modify the destination database that is being migrated.

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 migration, 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](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section describes how to use DRS to migrate data from a Redis Cluster instance to GeminiDB Redis over a public network.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 2-69 Migration task information

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

Region:

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

Project:

* Task Name: ⓘ

Description: ⓘ

0/256

Table 2-90 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 2-70 Replication instance information

Replication Instance Details ⓘ

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

• Data Flow: To the cloud Out of the cloud
The destination database must be a database in the current cloud. If you want to migrate data between databases, select To the cloud.

• Source DB Engine: MySQL MySQL schema and topic table MongoDB Redis

• Destination DB Engine:

• Network Type: ⓘ
DRS will automatically bind the specified EIP to the DRS instance and release the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.

• Destination DB Instance: ⓘ

• Replication Instance Subnet: ⓘ

• Migration Type: ⓘ
This migration type allows you to migrate data with minimal downtime. After a full migration finishes the destination database, an incremental migration pattern logs to ensure data consistency between the source and destination databases.

• DRS Task Nodes: ⓘ

• Specify EIP: ⓘ

Table 2-91 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.

Parameter	Description
Source DB Engine	Select Redis cluster .
Destination DB Engine	Select GeminiDB Redis .
Network Type	<p>Available options: VPC, VPN or Direct Connect, and Public network. Public network is selected in this example.</p> <ul style="list-style-type: none"> - VPC is suitable for migrations between cloud databases of the same account in the same region and VPC. - Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases. - VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	Select the DB instance you created.
Replication Instance Subnet	<p>The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Migration Type	<ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>

Parameter	Description
Source Shard Quantity	Number of subtasks connected to the source Redis Cluster instance. The value ranges from 1 to 16 and cannot be greater than the number of shards in the source Redis Cluster instance. Set this parameter based on the scale of the source Redis Cluster instance. You are advised to set one subtask to connect to four shards in the source Redis Cluster instance.
Specify EIP	This parameter is available when you select Public network for Network Type . Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances. For details about the data transfer fee generated using a public network, see EIP Price Calculator .

- AZ

Figure 2-71 AZ




Table 2-92 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 2-72 Enterprise projects and tags




Table 2-93 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 projects, 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 replication instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 2-73 DNS Server




Table 2-94 DNS server information

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

 **NOTE**

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

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

- Source database information

Figure 2-74 Source database information

Source Database

IP Address or Domain Name ?

Database Password ?

Table 2-95 Source database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the source database in the IP address/Domain name:Port format. Enter the IP addresses and ports of all shards in the source Redis Cluster instance. You are advised to enter the IP addresses of all shard slave nodes. Up to 32 IP addresses or domain names can be entered. Use commas (,) to separate multiple IP addresses or domain names. For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Password	Password used to log in to the source database.

 **NOTE**

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

- Destination database configuration

Figure 2-75 Destination database information

Destination Database

DB Instance Name


Database Password 

Table 2-96 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration task. This parameter cannot be changed.
Database Password	The password for the database username.

 **NOTE**

The password of the destination database is encrypted and stored in the database and the replication instance during the migration. After the task is deleted, the password is permanently deleted.

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

Figure 2-76 Migration object

Migration Object

Table 2-97 Migration object

Parameter	Description
Migration object	All database objects can be migrated.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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

Figure 2-77 Task startup settings




Table 2-98 Task startup settings

Parameter	Description
Start Time	Set Start Time to Start upon task creation or Start at a specified time based on site requirements. The Start at a specified time option is recommended. NOTE The migration task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.
Send Notifications	This parameter is optional. After enabled, select a SMN topic. If the task billing is about to start, the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.

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

----End

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

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration 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 Migration	Full+Incremental Migration
Source database user	<p>The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT</p> <p>If the source database version is 8.0, the user must have the SELECT permission for the mysql.user table. If the source database version is 5.7 or earlier, the user must have the SELECT permission for the MySQL system database.</p>	<p>The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <p>If the source database version is 8.0, the user must have the SELECT permission for the mysql.user table. If the source database version is 5.7 or earlier, the user must have the SELECT permission for the MySQL system database.</p>
Destination database user	<p>The user must have the following minimum permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</p> <p>To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database.</p>	

 **NOTE**

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, modify the connection information of the DRS task by referring to [Modifying Connection Information](#) to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.
- [Table 2-2](#) lists the minimum permissions required by a DRS task. If you need to migrate the grant permission through a DRS task, ensure that the connection account of the DRS task has the corresponding permission. Otherwise, the destination database user may not be authorized due to grant execution failure. For example, if the connection account of the DRS task does not require the process permission, but you need to migrate the process permission through a DRS task, ensure that the connection account of the DRS task has the process permission.

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 3-3](#). DRS will automatically check the objects you selected before the migration.

Table 3-3 Migration objects

Type	Precautions
Migration objects	<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: '<' '>' '^'• You can select table-, database-, or instance-level (all-level) migration.• Databases, tables, users, views, indexes, constraints, functions, stored procedures, triggers, and events can be migrated.• Only MyISAM and InnoDB tables can be migrated.• The system database and event statuses cannot be migrated.• Non-standard floating-point data that can be written in loose mode but cannot be written in strict mode cannot be migrated. Such non-standard floating-point data may fail to be hit, causing data migration failures.• Full and incremental migrations do not support invisible columns. Invisible columns can be migrated 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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-4](#).

Table 3-4 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database names cannot contain non-ASCII characters, or the following characters: '<code><>/\</code>'• The source table and view names cannot contain non-ASCII characters, or the following characters: '<code><>/\</code>'• The column names in the source database tables cannot end with a backslash (<code>\</code>).• The source database does not support the mysql binlog dump command.• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.• 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.• Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database.

Type	Restrictions
Restrictions on usage	<p data-bbox="587 297 695 327">General</p> <ul data-bbox="587 340 1430 1451" style="list-style-type: none"><li data-bbox="587 340 1430 472">● During the migration, 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.<li data-bbox="587 486 1430 584">● During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated.<li data-bbox="587 598 1430 663">● During migration, do not write the statement-based binlog into the source database.<li data-bbox="587 676 1430 741">● During migration, do not clear the binlog in the source database.<li data-bbox="587 754 1430 819">● During migration, do not create a database named ib_logfile in the source database.<li data-bbox="587 833 1430 994">● During the migration, 0 cannot be written to the auto-increment primary key column in the source database. Otherwise, the data of the auto-increment column in the source database is inconsistent with that in the destination database.<li data-bbox="587 1008 1430 1240">● During the migration, do not upgrade the source MySQL database across major versions. Otherwise, data may become inconsistent or the migration task may fail (data, table structures, and keywords may cause compatibility changes after the cross-version upgrade). You are advised to create a migration task again if the source MySQL database is upgraded across major versions.<li data-bbox="587 1254 1430 1451">● If index migration 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 data-bbox="587 1464 783 1494">Full migration</p> <ul data-bbox="587 1507 1430 1606" style="list-style-type: none"><li data-bbox="587 1507 1430 1606">● During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p data-bbox="587 1619 895 1648">Incremental migration</p> <ul data-bbox="587 1662 1430 1955" style="list-style-type: none"><li data-bbox="587 1662 1430 1727">● During an incremental migration of table-level objects, renaming tables is not supported.<li data-bbox="587 1740 1430 1805">● During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database.<li data-bbox="587 1818 1430 1955">● During an incremental migration, 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.

Type	Restrictions
	<ul style="list-style-type: none">• DDL statements are supported in the incremental migration phase. <p>Stopping a task</p> <ul style="list-style-type: none">• Stop a task normally. The selected events and triggers are migrated while the migration task proceeds to the final stage. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the migration status reported by the migration log.• Forcibly stop a task. If you forcibly stop a task, DRS resources will be released and triggers and events will not be migrated. You need to manually migrate triggers and events. If you want DRS to migrate triggers and events, restore the DRS task first. After the task status becomes normal, stop the task. <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 migration 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.• When creating multiple migration tasks in the many-to-one scenario, ensure that the read and write settings of the destination database are consistent in these tasks.• 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 a physically generated column in a table is generated based on a time type, the data in the column may be inconsistent.• The destination database cannot be restored to a point in time when a full migration was being performed.• If Transparent Data Encryption (TDE) is enabled for the source database table, TDE must also be enabled for the destination database. Otherwise, the table structure fails to be created and the task will be abnormal. You can create an unencrypted table structure in the destination database to avoid this problem.• If the destination MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.• If the destination database version is 5.7, the last digit 0 after the decimal point is lost in the floating point number of the JSON type due to version restrictions. The value comparison result will be inconsistent due to precision loss.• During the full and incremental migration phases, 0 cannot be written to the auto-increment primary key column in the source database. Otherwise, the data of the auto-increment column in the source database is inconsistent with that in the destination database.• 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.

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 migration, 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](#).
- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

This section uses the migration from an RDS for MySQL database to a MySQL database on an ECS as an example to describe how to configure a migration task in a VPC network on the DRS management console.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 3-1 Migration task information

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

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

Project

* Task Name ⓘ

Description ⓘ 0/256

Table 3-5 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 3-2 Replication instance information




Table 3-6 Replication instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> – VPC is suitable for migrations between cloud databases of the same account in the same region. – VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. – Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases.
Source DB Instance	Select the DB instance whose data is to be migrated out of the cloud.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects and data in non-system databases are migrated to the destination database at one time. The objects include tables, views, and stored procedures. <p>NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.</p> <ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-3 AZ




Table 3-7 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 3-4 Enterprise projects and tags

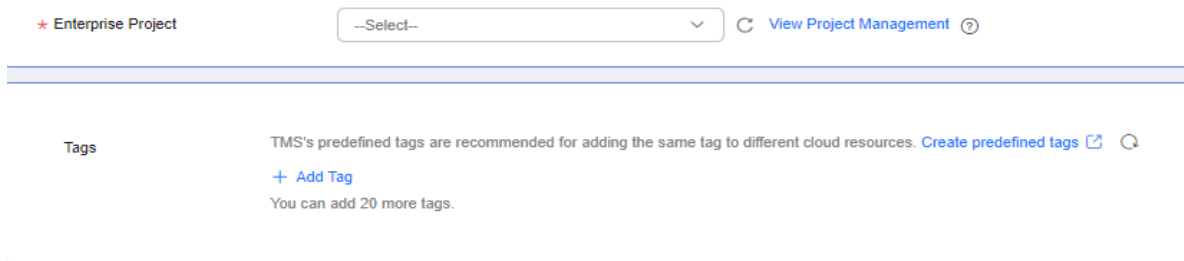


Table 3-8 Enterprise Project and Tags

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 projects, 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 replication instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 3-5 DNS Server

Configure Your Own DNS Server ⓘ

DNS Server

DNS Server IP Address . .

Table 3-9 DNS server information

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

NOTE

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

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

- Source database information

Figure 3-6 Source database information

Source Database

DB Instance Name

Database Username

Database Password

SSL Connection

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

Table 3-10 Source database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during migration task creation. This parameter cannot be changed.
Database Username	Enter the username of the source database.
Database Password	The password for the database username. If the task is in the Starting, Full migration, Incremental migration, or Incremental migration 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 username and password of the source database are encrypted and stored in the database and the replication instance during the migration. After the task is deleted, the username and password are permanently deleted.

- Destination database information

Figure 3-7 Destination database information

Destination Database

VPC [View VPC](#)

Subnet [View Subnets](#)

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Migrate Definer to User Yes No

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

Table 3-11 Destination database settings

Parameter	Description
VPC	A dedicated virtual network in which the destination database is located. It isolates networks for different services.
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	Enter 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: If the task is in the Starting, Full migration, Incremental migration, or Incremental migration 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.

Parameter	Description
Migrate Definer to User	<p>Indicates whether to migrate the Definers of all source database objects to the destination database user entered during the connection test.</p> <ul style="list-style-type: none"> - Yes The Definers of all source database objects will be migrated to the user. Other users do not have permissions for database objects unless these users are authorized. For details about authorization, see How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration? <p>For example, if the view is <code>CREATE ALGORITHM=UNDEFINED DEFINER='username'@'%` SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> before migration,</p> <p>it is converted to <code>CREATE ALGORITHM=UNDEFINED DEFINER='drsUser'@'%` SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> after the migration.</p> <p>drsUser indicates the destination database user used for testing the connection.</p> <ul style="list-style-type: none"> - No The Definers of all source database objects will not be changed. You need to migrate all accounts and permissions of the source database in the next step. Note that if the Definer account is not found in the destination database, unavailable objects will be created. <p>For details about Definer, see the MySQL official document.</p>

 **NOTE**

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

Step 4 On the **Set Task** page, set migration accounts and objects, and click **Next**.

Figure 3-8 Migration Type

Flow Control Yes No

Migrate Incremental Accounts and Permissions Yes No
The migration of incremental accounts and permissions may fail because the source and destination MySQL versions and account encryption modes may be different.

Migrate Account Yes No
During a database migration, you need to separately migrate accounts and permissions. Certain accounts cannot be migrated to the destination database. Ensure that services are not affected.

<input checked="" type="checkbox"/> Account	Account	Can Be Migrated	Permission	Password	Remarks
<input checked="" type="checkbox"/>	root@% %	Yes	GRANT SELECT INSERT UPDATE DELETE CREATE DROP REL..	<input type="text" value=""/>	--
<input checked="" type="checkbox"/>	root@% %	Yes	GRANT SELECT INSERT UPDATE DELETE ON **	<input type="text" value=""/>	--
<input checked="" type="checkbox"/>	root@% %	Yes	GRANT USAGE ON **	<input type="text" value=""/>	--
<input checked="" type="checkbox"/>	mysql@% %	Yes	GRANT SELECT INSERT UPDATE DELETE CREATE DROP REL..	<input type="text" value=""/>	--
<input checked="" type="checkbox"/>	mysql@% %	Yes	GRANT USAGE ON **	<input type="text" value=""/>	--
<input checked="" type="checkbox"/>	user_mvc@% %	Yes	GRANT SELECT ON **	<input type="text" value=""/>	--

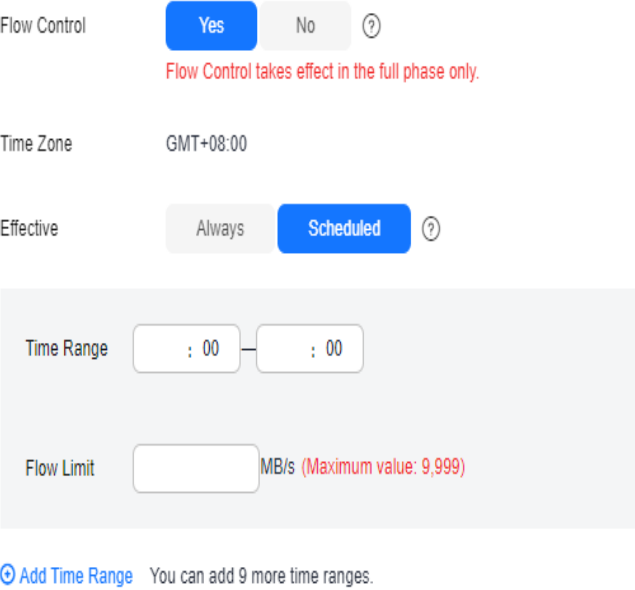
Reset Password
 Set Unlocked Password

Create Indexes Along With Table Structure Yes No


Filter DROP DATABASE Yes No

Migration Object All Tables Databases

Table 3-12 Migration types and objects

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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-9 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> - Flow control mode takes effect only during a full migration. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Incremental Accounts and Permissions	<p>Indicates whether to migrate incremental accounts and permissions during database migration.</p> <ul style="list-style-type: none"> ● Yes All incremental accounts and permissions will be migrated. The migration of incremental accounts and permissions may fail because the source and destination database versions and account encryption modes may be different. ● No All incremental accounts and permissions will be filtered out during the migration.
Filter DROP DATABASE	<p>During an incremental migration, executing DDL operations on the source database may affect the data migration performance to some extent. To reduce data migration risks, DRS allows you to filter out DDL operations.</p> <p>The database deletion operation can be filtered out by default.</p> <ul style="list-style-type: none"> ● If you select Yes, any database deletion operations performed on the source database are not synchronized during data migration. ● If you select No, related operations are synchronized to the destination database during data migration.
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none"> ● Yes If you need to migrate accounts, see Migrating Accounts. ● No During the migration, accounts, permissions, and passwords are not migrated.
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full migration phase.</p> <ul style="list-style-type: none"> ● Yes: Indexes are migrated when the table structure is migrated in the full migration phase. ● No: Indexes are migrated separately after data migration.

Parameter	Description
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> • All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. • Tables: The selected table-level objects will be migrated. • Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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

Figure 3-10 Task startup settings

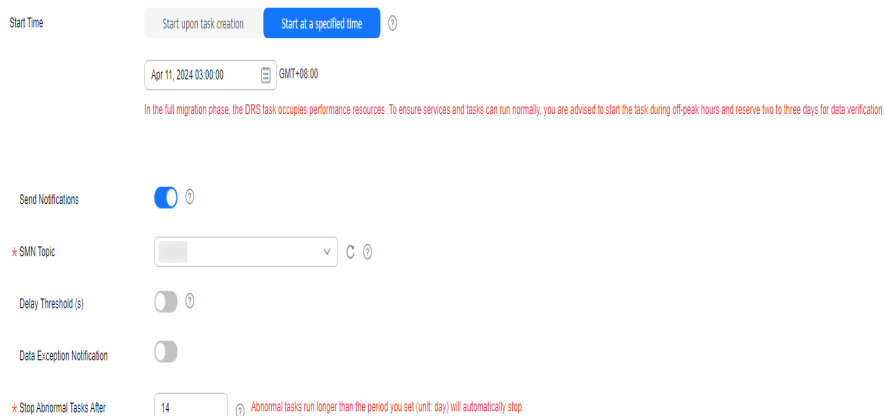



Table 3-13 Task startup settings

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

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.2 From DDS to MongoDB

Supported Source and Destination Databases

Table 3-14 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> • DDS DB instances (versions 3.2, 3.4, 4.0, 4.2, 4.4 and 5.0) <p>NOTE DDS 5.0 supports replica sets only.</p>	<ul style="list-style-type: none"> • On-premises MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) • ECS-hosted MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) • Other cloud MongoDB (versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4 and 5.0) <p>NOTE The destination database version must be the same as or later than the source database version. DDS 5.0 supports replica sets only.</p>

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 3-15](#). DRS will automatically check the objects you selected before the migration.

Table 3-15 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none"> • You can select table-, database-, or instance-level (all-level) migration. • Replica set: Only collections (including validator and capped collections), indexes, and views can be migrated. • Cluster: Only collections (including validator and capped collections), shard keys, indexes, and views can be migrated. • Single node: Only collections (including validator and capped collections), indexes, and views can be migrated. • System databases (such as local, admin, and config) cannot be migrated. You need to manually create usernames and roles in the destination database. • The statement for creating a view cannot contain a regular expression. • Collections that contain the <code>_id</code> field without indexes are not supported. • The first parameter of <code>BinData()</code> cannot be 2. • Do not store non-UTF-8 character strings in the String field of the source database collection. Otherwise, data will be inconsistent before and after the migration. • Time series collections are not supported in both the full and incremental phases.

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 3-16](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

Table 3-16 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	<ul style="list-style-type: none"> • Replica set: The source database user must have the readAnyDatabase permission for the admin database. • Single node: The source database user must have the readAnyDatabase permission for the admin database. • Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. 	<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.
Destination database user	<p>The user who connects to the destination database must have the dbAdminAnyDatabase permission of the admin database and the readWrite permission of the destination database.</p> <p>If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.</p>	

 **NOTE**

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

```
db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"},{role:"read",db:"config"}])
```

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 [DRS Pre-Check Items](#). In addition to the pre-check items, you need to pay attention to the items listed in [Table 3-17](#).

Table 3-17 Precautions

Type	Restrictions
Restrictions on the source database	<ul style="list-style-type: none">• The source database name cannot contain \. \$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).• When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique.

Type	Restrictions
Use restrictions	<p>General</p> <ul style="list-style-type: none"> ● During the migration, do not modify or delete the usernames, passwords, and permissions of the source and destination databases or change the ports of the source and destination databases. ● During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. ● During migration, data rollback caused by a primary/standby switchover of the source database is not supported. ● During the migration, documents larger than 16 MB in the source database cannot be inserted or updated. ● During the migration, do not run sh.moveChunk() on the source database. Otherwise, the migrated data will be inconsistent. <p>Full migration</p> <ul style="list-style-type: none"> ● During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. <p>Incremental migration</p> <ul style="list-style-type: none"> ● During an incremental migration of collections, you are advised not to rename the collections. ● To ensure the migration performance, concurrent replay is performed at the set level in the incremental task phase. In the following scenarios, only single-thread write is supported and concurrent replay is not supported: <ul style="list-style-type: none"> - The collection index contains a unique key. - The value of capped of the collection attribute is true. <p>In either of the preceding scenarios, the task delay may increase.</p> <p>Data comparison</p> <ul style="list-style-type: none"> ● You are advised to compare data during off-peak hours of the source database to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks. ● During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent. ● During row comparison, if an orphan document exists in a cluster instance or chunks are being migrated, the number of returned rows is incorrect and the comparison results are inconsistent. For details, see MongoDB official documents.

Type	Restrictions
Other restrictions	<ul style="list-style-type: none"> ● If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental migration, data consistency cannot be ensured when source and destination databases are in different time zones or clocks. ● 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 migrated. 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 migrate compression parameters because the destination database does not support compression parameter settings. ● If the destination database is on 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 migration task. If you enter information about primary and secondary nodes, ensure that all nodes belong to the same replica set instance. ● If the destination database is on a cluster instance out of the cloud, you are advised to enter information about multiple mongos nodes to reduce the impact of the single-node failure on the migration task. In addition, multiple mongos nodes support load balancing. In addition, ensure that all mongos nodes belong to the same cluster instance. ● To accelerate the migration, delete unnecessary indexes from the source database and retain only necessary indexes before the migration. You are advised not to create indexes for the source database during the migration. 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.

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 migration, 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](#).

- You have read [DRS Pre-Check Items](#) and [Precautions](#).

Procedure

This section describes how to migrate data from a DDS instance to an on-premises MongoDB database over a public network.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 3-11 Migration task information

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

Region: [Region] ...

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

Project: [Project] ...

* Task Name: DRS-5678 ⓘ

Description: [Description] ⓘ

0/256

Table 3-18 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 3-12 Replication instance information

Replication Instance Details ⓘ

The following information cannot be modified after you click the **Create Now** button.

• Data Flow: To the cloud [Out of the cloud](#)

Out of the cloud: The source database must be a database in the current cloud.

• Source DB Engine: MySQL [GeneralDB Proxy](#) [DDS](#)

• Destination DB Engine: [MySQL](#)

• Network Type: Public network ⓘ

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

• Source DB Instance: No DB instance available ⓘ [View DB Instance](#) [View Unavailable DB Instance](#)

• Replication Instance Subnet: Search the subnet ⓘ [View Subnet](#) [View Occupied IP Address](#)

• Migration Type: Full-Instance [Full](#) ⓘ

The migration type allows you to migrate data with minimal downtime. After a full migration releases the destination database, an incremental migration saves logs to ensure data consistency between the source and destination databases.

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

Table 3-19 Replication instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is an RDS database on the current cloud or a DDS DB instance. It is required that either the source database or the destination database is on the current cloud.
Source DB Engine	Select DDS .
Destination DB Engine	Select MongoDB .
Network Type	Available options: Public network, VPC, VPN or Direct Connect <ul style="list-style-type: none"> - VPC is suitable for migrations between cloud databases of the same account in the same region. - VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. - Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases.
Source DB Instance	Select the DDS DB instance to be migrated.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is permitted. It migrates all objects and data in non-system databases to the destination database at one time. The objects include collections and indexes. <p>NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.</p> <ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-13 AZ




Table 3-20 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 3-14 Enterprise projects and tags

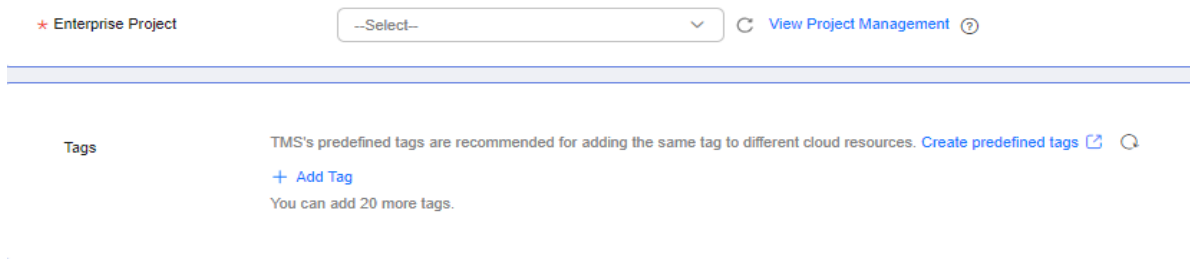


Table 3-21 Enterprise Project and Tags

Parameter	Description
Enterprise Project	<p>An enterprise project you would like to use to centrally manage your cloud resources and members. Select an enterprise project from the drop-down list. The default project is default.</p> <p>For more information about enterprise projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Figure 3-15 Source database information

Destination Database

DB Instance Name

Authentication Database

Database Username

Database Password

SSL Connection

Table 3-22 Source database information

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration 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	Enter the username of the source database.
Database Password	Enter the password of the source database user.
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 the database and the replication instance during the migration. After the task is deleted, the username and password are permanently deleted.

Figure 3-16 Destination database information

Destination Database

IP Address or Domain Name ?
Ensure that the entered addresses belong to the same DB instance.

Authentication Database

Database Username

Database Password

SSL Connection

Table 3-23 Destination database settings

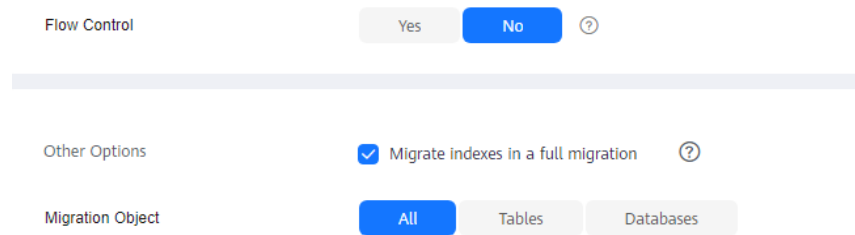
Parameter	Description
IP Address or Domain Name	<p>IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535</p> <p>You can enter a maximum of three groups of IP addresses or domain names of the destination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

 **NOTE**

The IP address, domain name, 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 Task** page, select migration objects and click **Next**.

Figure 3-17 Migration object



Flow Control

Yes No ?

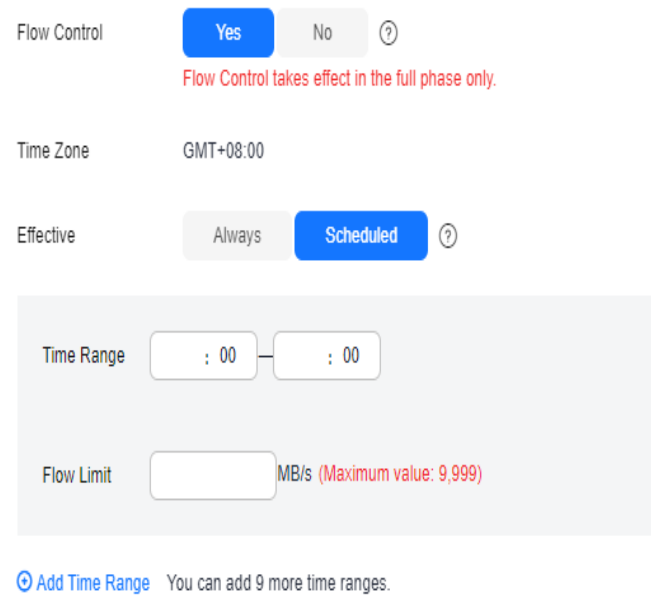
Other Options


Migrate indexes in a full migration ?

Migration Object

All Tables Databases

Table 3-24 Migrate 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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 3-18 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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"> - Flow control mode takes effect only during a full migration. - You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Other Options	Determine whether to migrate the indexes you create during full migration. The default index based on <code>_id</code> is automatically created in the destination. If indexes are not migrated, the indexes are not compared.
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> • All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. • Tables: The selected table-level objects will be migrated. • Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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

Figure 3-19 Task startup settings

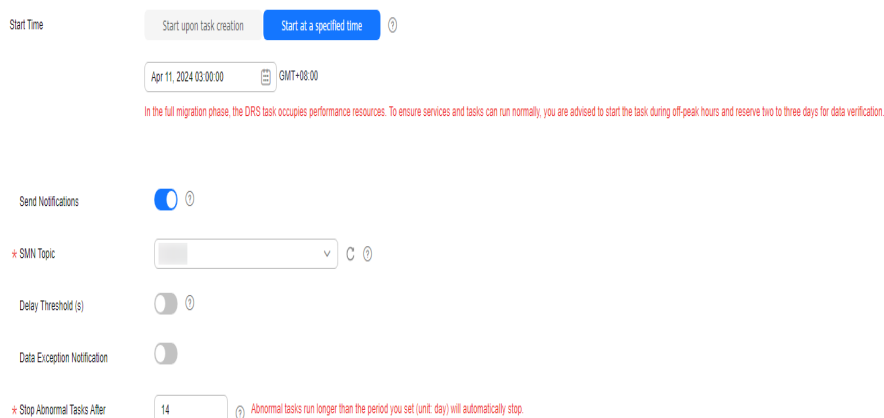



Table 3-25 Task startup settings

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

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.3 From GeminiDB Redis to Redis

Supported Source and Destination Databases

Table 3-26 Supported databases

Source DB	Destination DB
GeminiDB Redis	<ul style="list-style-type: none">• On-premises Codis clusters (open-source Codis 3.0 or later)• Codis clusters on ECSs (open-source Codis 3.0 or later)• On-premises single-node Redis 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x• ECS-hosted single-node Redis 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x• On-premises master/standby Redis 4.0.x and 5.0.x• ECS-hosted master/standby Redis 4.0.x and 5.0.x• GeminiDB Redis

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 3-27](#). DRS will automatically check the objects you selected before the migration.

Table 3-27 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">• Object level: all database objects• Supported migration objects:<ul style="list-style-type: none">– Data of the String, Hash, List, Set, and Sorted Set types can be migrated.– Lua scripts and transactions cannot be migrated.– Stream, Exhash, and Bloomfilter data types cannot be migrated.– User-defined types are not supported.– Commands that are not supported by the destination database cannot be migrated.

Suggestions

⚠ CAUTION

To maintain data consistency before and after the migration, ensure that no data is written to your source and destination databases during a full migration. In the full+incremental migration mode, you can continue the migration while data is still being written to the source database.

- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
- In the migration, ensure that no data is written to the destination database to ensure data consistency before and after the migration.
- Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full migration, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental migration 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 migration, read the following precautions before creating a migration task.

Table 3-28 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> - The source must be a GeminiDB Redis instance on the current cloud. - A GeminiDB Redis instance can be used as the source database of only one DRS task. - The source GeminiDB Redis instance kernel version (gemini version) must be 4.2.0 or later. - The sequence number of the source database with data stored must be less than the maximum sequence number of the destination database. For example, if the source database contains No.10 database with data stored, the maximum sequence number of the destination database must be greater than or equal to 10. - To create a full+incremental task, the PSYNC command must be supported. - The replication function must be enabled for the source GeminiDB Redis instance. <code>set config set enable-replication 1</code> - To use data migration, disable the configuration item notify-keyspace-events. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database must be empty before the migration. - The destination DB instance must have sufficient storage space. - If the destination is master/standby Redis 5.0, enter all IP addresses and ports of the master, slave, and sentinel nodes for the destination database connection testing. The passwords of the master, slave, and sentinel nodes must be the same. - If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is set on the master and slave nodes, DRS regards the master/standby Redis as a single-node Redis. The master/standby switchover of the destination database is not supported. For the destination database connection testing, enter only the IP address and port of the master node. - If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is not set on the master or slave node, the master/standby switchover of the destination database is supported. For the destination database connection testing, enter all IP addresses and ports of the master, slave, and sentinel nodes. ● Other notes: <ul style="list-style-type: none"> - A full+incremental task read data by parsing RDB files.

Type	Restrictions
	<ul style="list-style-type: none"> - If the replication timeout interval of the source GeminiDB Redis is too short (the default value is 3600s), you can increase the value of repl-timeout to prevent task resuming failure due to a long task failure. - During the migration, if the task fails due to configuration increase of the source GeminiDB Redis, you need to reset the task. To ensure data consistency, clear the data that has been migrated to the destination database before reconfiguring the task. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.) - During the migration, the specifications of the source GeminiDB Redis cannot be downgraded. If the specifications of the source database needs to be downgraded, create a migration task again. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.) - For List objects, duplicated data may occur because the existing data on the destination end is not cleared when PSYNC is invoked for retransmission. - If the source is a primary/standby GeminiDB Redis instance, commands containing multiple keys (such as mset) cannot be executed in the source database. Otherwise, the error message "CROSSSLOT Keys in request don't hash to the same slot" may be displayed because the keys are not in the same slot or involve multiple slots.
Full migration	<ul style="list-style-type: none"> • During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases. • During the migration, do not modify the destination database that is being migrated.
Incremental migration	<ul style="list-style-type: none"> • During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases. • During the migration, do not modify the destination database that is being migrated.

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 migration, 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](#).

- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section describes how to use DRS to migrate data from GeminiDB Redis to a Redis database through a VPN.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 3-20 Migration task information

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

Region

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

Project

* Task Name ⓘ

Description ⓘ

0/256

Table 3-29 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 3-21 Replication instance information



Table 3-30 Replication instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select GeminIDB Redis .
Destination DB Engine	Select Redis .
Network Type	Available options: VPC , VPN or Direct Connect , and Public network . VPN or Direct Connect is selected in this example. <ul style="list-style-type: none"> – VPC is suitable for migrations between cloud databases of the same account in the same region and VPC. – Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases. – VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	Select the source GeminIDB Redis instance you have created.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
DRS Task Nodes	Specifies the number of nodes to be migrated using DRS. One DRS task node connects to two to three shards on the source database. If the source GeminiDB Redis database is a primary/standby instance, the number of DRS task nodes can only be 1.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. The number of specified EIPs must be the consistent with that of DB instances.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 3-22 AZ




Table 3-31 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 3-23 Enterprise projects and tags




Table 3-32 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 projects, 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 replication instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 3-24 DNS Server

Configure Your Own DNS Server ⓘ

DNS Server

DNS Server IP Address

Table 3-33 DNS server information

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

NOTE

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

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

- Source database information

Figure 3-25 Source database information

Source Database

DB Instance Name

Database Password

Table 3-34 Source database settings

Parameter	Description
DB Instance Name	The GeminiDB Redis instance selected when you created the migration task. This parameter cannot be changed.
Database Password	Password of the source database administrator.


 **NOTE**

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

- Destination database configuration

Figure 3-26 Destination database information

Destination Database

IP Address or Domain Name 

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

Database Password

Table 3-35 Destination database settings

Parameter	Description
IP Address or Domain Name	<p>IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535</p> <ul style="list-style-type: none"> – If a single-node or proxy instance is used, ensure that the entered IP addresses or domain names belong to the same instance. – If the destination is master/standby Redis 5.0, enter all IP addresses and ports of the master, slave, and sentinel nodes for the destination database connection testing. The passwords of the master, slave, and sentinel nodes must be the same. – If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is set on the master and slave nodes, DRS regards the master/standby Redis as a single-node Redis. The master/standby switchover of the destination database is not supported. For the destination database connection testing, enter only the IP address and port of the master node. – If the destination is master/standby Redis 4.0 and the requirepass parameter (a password) is not set on the master or slave node, the master/standby switchover of the destination database is supported. For the destination database connection testing, enter all IP addresses and ports of the master, slave, and sentinel nodes.
Database Password	Password for logging in to the destination database

 **NOTE**

The password of the destination database is encrypted and stored in the database and the replication instance during the migration. After the task is deleted, the password is permanently deleted.

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

Figure 3-27 Migration object



Table 3-36 Migration object

Parameter	Description
Migration Object	All database objects can be migrated.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

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


Figure 3-28 Task startup settings



Table 3-37 Task startup settings

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

3.4 From GeminiDB Redis to Redis Cluster

Supported Source and Destination Databases

Table 3-38 Supported databases

Source DB	Destination DB
GeminiDB Redis	<ul style="list-style-type: none">• ECS-hosted Redis Cluster (open-source Redis 4.0 or later)• On-premises Redis Cluster (open-source Redis 4.0 or later)

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 3-39](#). DRS will automatically check the objects you selected before the migration.

Table 3-39 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">• Object level: all database objects• Supported migration objects:<ul style="list-style-type: none">– Data of the String, Hash, List, Set, and Sorted Set types can be migrated.– Lua scripts and transactions cannot be migrated.– Stream, Exhash, and Bloomfilter data types cannot be migrated.– User-defined types are not supported.– Commands that are not supported by the destination database cannot be migrated.– Batch operation commands without hash_tag cannot be migrated.

Suggestions

⚠ CAUTION

To maintain data consistency before and after the migration, ensure that no data is written to your source and destination databases during a full migration. In the full+incremental migration mode, you can continue the migration while data is still being written to the source database.

-
- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
 - In the migration, ensure that no data is written to the destination database to ensure data consistency before and after the migration.
 - Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full migration, and two to four CPUs are occupied.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - For more information about the impact of DRS on databases, see [How Does DRS Affect the Source and Destination Databases?](#)

Precautions

The full+incremental migration 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 migration, read the following precautions before creating a migration task.

Table 3-40 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database requirements: <ul style="list-style-type: none"> - The source must be a GeminiDB Redis instance on the current cloud. - A GeminiDB Redis instance can be used as the source database of only one DRS task. - The source GeminiDB Redis instance kernel version (gemini version) must be 4.2.0 or later. - To use data migration, disable the configuration item notify-keyspace-events. - The replication function must be enabled for the source GeminiDB Redis instance. <code>set config set enable-replication 1</code> - To create a full+incremental task, the PSYNC command must be supported. ● Destination database requirements: <ul style="list-style-type: none"> - The destination database must be empty before the migration. - The destination DB instance must have sufficient storage space. - The username cannot be entered in Redis 6.0. - If the destination cluster instance is created using containers, ECSs, or NAT, ensure that the config settings of the cluster allow DRS to discover all master and standby nodes. ● Other notes: <ul style="list-style-type: none"> - A full+incremental task read data by parsing RDB files. - If the replication timeout interval of the source GeminiDB Redis is too short (the default value is 3600s), you can increase the value of repl-timeout to prevent task resuming failure due to a long task failure. - During the migration, if the task fails due to configuration increase of the source GeminiDB Redis, you need to reset the task. To ensure data consistency, clear the data that has been migrated to the destination database before reconfiguring the task. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.) - During the migration, the specifications of the source GeminiDB Redis cannot be downgraded. If the specifications of the source database needs to be downgraded, create a migration task again. (You do not need to reset the task if the task is in the incremental state and the source database kernel version (Gemini version) is 6.3.0 or later.)

Type	Restrictions
	<ul style="list-style-type: none"> - For List objects, duplicated data may occur because the existing data on the destination end is not cleared when PSYNC is invoked for retransmission. - If the source is a primary/standby GeminiDB Redis instance, commands containing multiple keys (such as mset) cannot be executed in the source database. Otherwise, the error message "CROSSSLOT Keys in request don't hash to the same slot" may be displayed because the keys are not in the same slot or involve multiple slots.
Full migration	<ul style="list-style-type: none"> • During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases. • During the migration, do not modify the destination database that is being migrated.
Incremental migration	<ul style="list-style-type: none"> • During migration, do not change the passwords of the source and destination databases or change the ports of the source and destination databases. • During the migration, do not modify the destination database that is being migrated.

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 migration, 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](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section describes how to use DRS to migrate data from GeminiDB Redis to a Redis Cluster instance through a VPN.

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 3-29 Migration task information

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

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

Project:

* Task Name: ⓘ

Description: ⓘ
0/256

Table 3-41 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 3-30 Replication instance information

Replication Instance Details ⓘ

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

- Data Flow: To the cloud **Out of the cloud**
- Source DB Engine:
- Destination DB Engine:
- Network Type: ⓘ
- Source DB Instance: ⓘ New DB Instance Use Unavailable DB Instance
- Replication Instance Subnet: ⓘ
- Migration Type:
This migration type allows you to migrate data with minimal downtime. After a full migration replicates the destination database, an incremental migration process logs to ensure data consistency between the source and destination databases.
- DRS Task Nodes: ⓘ
- Specify EP: ⓘ

Table 3-42 Replication instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select GeminiDB Redis .

Parameter	Description
Destination DB Engine	Select Redis cluster .
Network Type	Available options: VPC , VPN or Direct Connect , and Public network . VPN or Direct Connect is selected in this example. <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region and VPC.- Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases.- VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Source DB Instance	Select the source GeminiDB Redis instance you have created.
Replication Instance Subnet	The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.
Migration Type	<ul style="list-style-type: none">- Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and 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 migration will be continuously synchronized to the destination database, and the source remains accessible.</p>
DRS Task Nodes	Specifies the number of nodes to be migrated using DRS. One DRS task node connects to two to three shards on the source database. If the source GeminiDB Redis database is a primary/standby instance, the number of DRS task nodes can only be 1.

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

- AZ

Figure 3-31 AZ




Table 3-43 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 3-32 Enterprise projects and tags




Table 3-44 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 projects, 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 replication instance is created. Then, configure your own DNS server, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 3-33 DNS Server




Table 3-45 DNS server information

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

 **NOTE**

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

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

- Source database information

Figure 3-34 Source database information

Source Database

DB Instance Name

Database Password

Table 3-46 Source database settings

Parameter	Description
DB Instance Name	The GeminiDB Redis instance selected when you created the migration task. This parameter cannot be changed.
Database Password	Password of the source database administrator.

 **NOTE**

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

- Destination database configuration

Figure 3-35 Destination database information

Destination Database

IP Address or Domain Name ⓘ

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

Database Password ⓘ

Table 3-47 Destination database settings

Parameter	Description
IP Address or Domain Name	IP address or domain name of the destination database in the IP address/Domain name:Port format. Enter the IP addresses and ports of all master and slave nodes in the Redis Cluster instance. A maximum of 32 IP addresses or domain names can be entered. Use commas (,) to separate multiple values. For example: 192.168.0.1:8080,192.168.0.2:8080.
Database Password	Password for logging in to the destination database.

 **NOTE**

The password of the destination database is encrypted and stored in the database and the replication instance during the migration. After the task is deleted, the password is permanently deleted.

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

Figure 3-36 Migration object



Table 3-48 Migration object

Parameter	Description
Migration Object	All database objects can be migrated.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

NOTE

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

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

Figure 3-37 Task startup settings

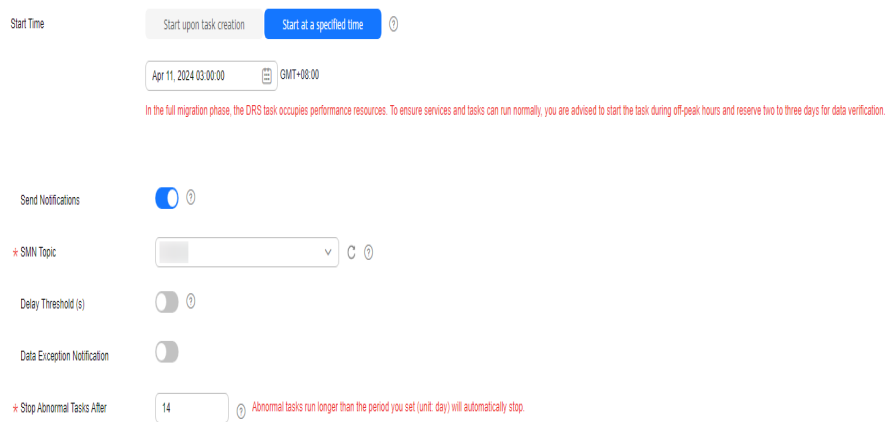



Table 3-49 Task startup settings

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

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

Step 7 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

4 Task Management

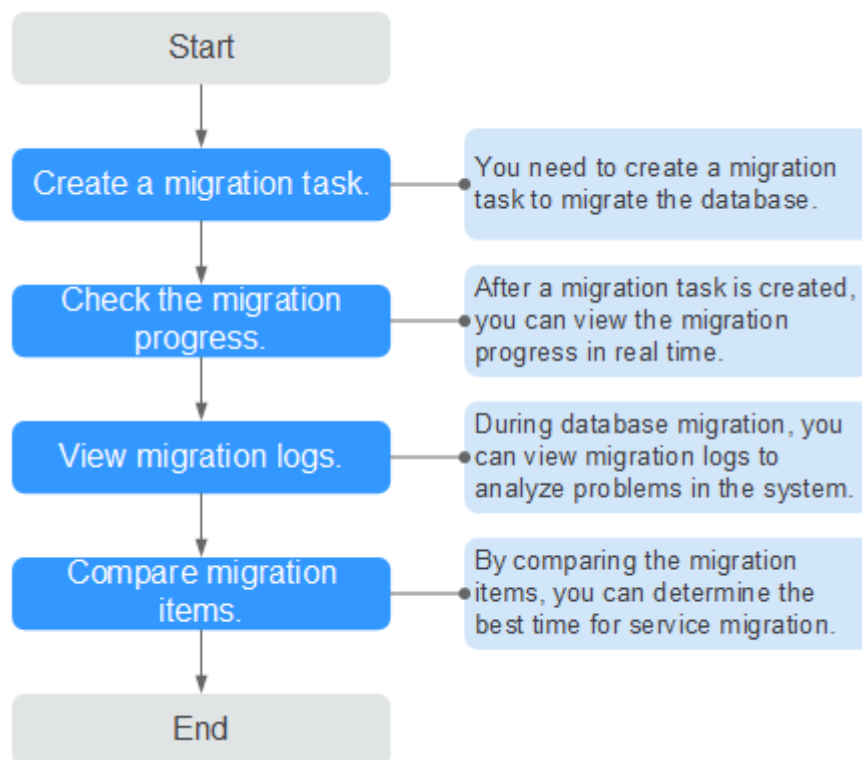
4.1 Creating a Migration Task

Process

A complete real-time migration consists of creating a migration task, tracking task progress, analyzing migration logs, and comparing data consistency. By comparing multiple items and data, you can determine the proper time for service migration to minimize the service downtime.

A complete migration involves the following procedures.

Figure 4-1 Migration process



- **Step 1: Create a migration task.** Select the source and destination databases as required and create a migration task.
- **Step 2: Check the migration progress.** During migration, you can view the migration progress.
- **Step 3: View migration logs.** Migration logs contain alarms, errors, and prompt information. You can analyze system problems based on such information.
- **Step 4: Compare migration items.** You can compare objects and data to be migrated to ensure data consistency.

This section uses the migration from MySQL to RDS for MySQL as an example to describe how to configure a migration task over a VPC network on the DRS console.

VPC is suitable for migrations of cloud databases in the same region.

You can create a migration task that will walk you through each step of the process. After a migration task is created, you can manage it on the DRS console.

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 migration, 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

Step 1 On the **Online Migration Management** page, click **Create Migration Task**.

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

- Task information description

Figure 4-2 Migration 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	<input type="text" value=".."/>	<small>Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.</small>
Project	<input type="text" value=".."/>	
* Task Name	<input type="text" value="DRS-5678"/>	<small>ⓘ</small>
Description	<input style="height: 40px;" type="text"/>	

0256

Table 4-1 Task information

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your services.
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 !=<>'&"\

- Replication instance information

Figure 4-3 Replication instance information

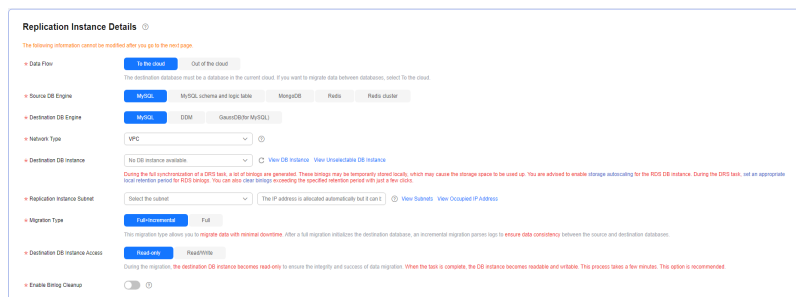


Table 4-2 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination DB is on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .

Parameter	Description
Network Type	<p>Select VPC Network.</p> <p>Available options: VPC, VPN or Direct Connect, and Public network. By default, the value is Public network.</p> <ul style="list-style-type: none"> - VPC is suitable for migrations between cloud databases of the same account in the same region and VPC. - Public network is suitable for migrations from on-premises databases or external cloud databases to destination databases. - VPN or Direct Connect is suitable for migrations from on-premises databases to cloud databases or between databases across regions on the cloud using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection. <p>For details about networks, see Preparations.</p>
Destination DB Instance	<p>The RDS DB instance you created.</p>
Replication Instance Subnet	<p>The subnet where the replication instance resides. 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 replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>

Parameter	Description
Destination DB Instance Access	<ul style="list-style-type: none"> <li data-bbox="691 297 1442 1003"> <p>– Read-only</p> <p>During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.</p> <p>If a migration task fails, the destination database restores to the read/write state. After the migration task is resumed, the destination database becomes read-only.</p> <p>The access settings of the destination instance are also affected by the access settings of the DB instance itself. Therefore, you are advised not to set the access settings of the DB instance on the RDS console.</p> <p>If the destination database is set to read-only for a migration task, you can disable read-only on the task details page. After read-only is disabled, the destination database cannot be set back to read-only. Even if the migration task fails and then resumed and started, the destination database cannot be set to read-only.</p> <li data-bbox="691 1010 1442 1444"> <p>– Read/Write</p> <p>During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks. Set the destination database to Read/Write only when you need to modify other data in the database during the migration.</p> <p>The task cannot be modified after being created.</p>

Parameter	Description
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects and data in non-system databases are migrated to the destination database at one time. The objects include tables, views, and stored procedures. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database. - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.
Enable Binlog Cleanup	Indicates whether to enable the function of quickly clearing binlogs of the destination database. After this function is enabled, binlog clearing is enabled for the destination database during the full migration and disabled during the incremental migration.
Specify EIP	<p>This parameter is available when you select Public network for Network Type. Select an EIP to be bound to the DRS instance. DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete.</p> <p>For details about the data transfer fee generated using a public network, see EIP Price Calculator.</p>

- AZ

Figure 4-4 AZ




Table 4-3 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 4-5 Enterprise projects and tags

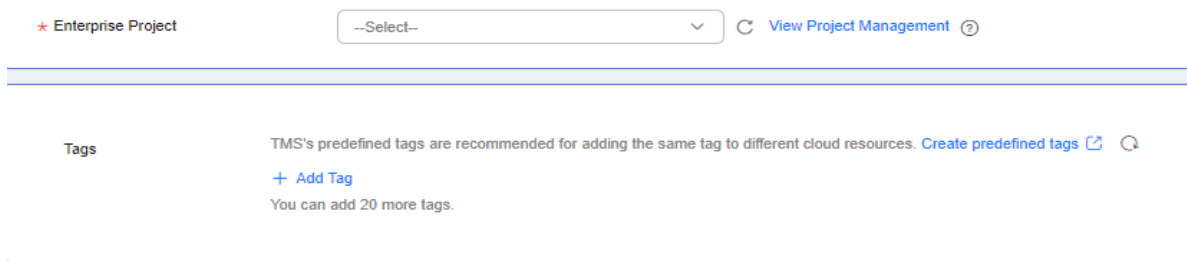


Table 4-4 Enterprise Project and Tags

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 projects, 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 replication 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 replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- (Optional) Configuring your own DNS server

Figure 4-6 DNS Server

Configure Your Own DNS Server 

DNS Server

DNS Server IP Address . .

Table 4-5 DNS server information

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

 **NOTE**

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

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

- Scenario 1: Self-built database - source database configuration

Figure 4-7 Self-built database - 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 Username:

Database Password:

SSL Connection:

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

Table 4-6 Self-built database - source database information

Parameter	Description
Database Type	Select Self-built database .
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 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 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 4-8 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.

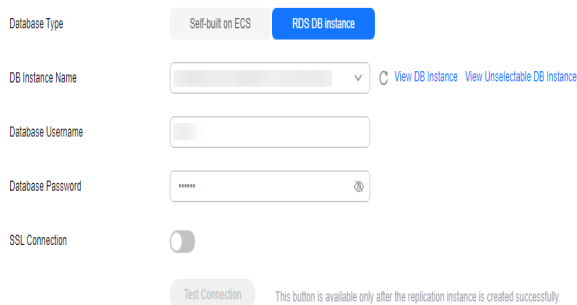


Table 4-7 RDS DB instance - source database information

Parameter	Description
Database Type	Select RDS DB instance .
DB Instance Name	Select the RDS DB instance to be migrated as the source DB instance.
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.

- Destination database configuration

Figure 4-9 Destination database information

Destination Database

DB Instance Name

Database Username

Database Password

Migrate Definer to User Yes No

SSL Connection

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

Table 4-8 Destination database settings

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

Parameter	Description
Database Password	The password for the database username.
Migrate Definer to User	<p>Indicates whether to migrate the Definers of all source database objects to the destination database user entered during the connection test.</p> <ul style="list-style-type: none">- Yes The Definers of all source database objects will be migrated to the user. Other users do not have permissions for database objects unless these users are authorized. For details about authorization, see How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration? <p>For example, if the view is <code>CREATE ALGORITHM=UNDEFINED DEFINER='username'@'%` SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> before migration,</p> <p>it is converted to <code>CREATE ALGORITHM=UNDEFINED DEFINER='drsUser'@'%` SQL SECURITY DEFINER VIEW `test_db`.`view5` AS select 1 AS `1`</code> after the migration.</p> <p>drsUser indicates the destination database user used for testing the connection.</p> <ul style="list-style-type: none">- No The Definers of all source database objects will not be changed. You need to migrate all accounts and permissions of the source database in the next step. Note that if the Definer account is not found in the destination database, unavailable objects will be created. <p>For details about Definer, see the MySQL official document.</p>
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 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 Task** page, select the accounts and objects to be migrated, and click **Next**.

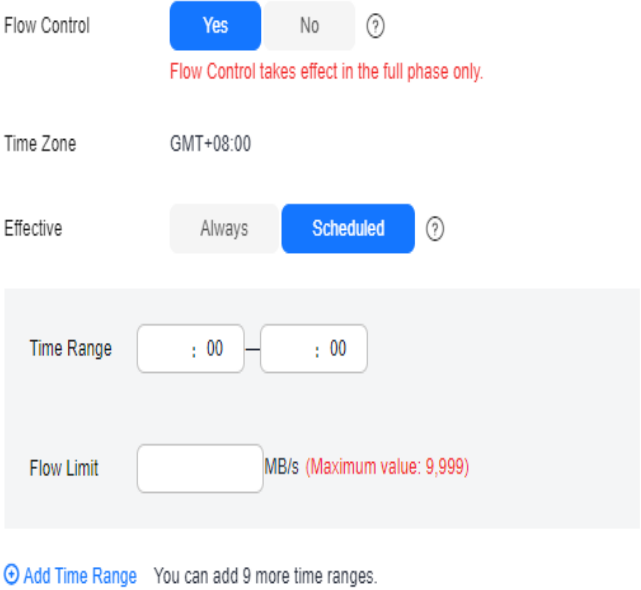
Figure 4-10 Migration type

The screenshot shows the 'Migrate account' section of the 'Set Task' page. It includes a table with the following columns: Account, Can be Migrated, Permissions, Password, and Remarks. The table lists several accounts and their migration status.


Account	Can be Migrated	Permissions	Password	Remarks
<input checked="" type="checkbox"/> <code>hadoop_hdfs@%*</code>	Yes	GRANT REPLICATION ADMIN, REPLICATION CLIENT ON *.* GRANT SELECT		
<input checked="" type="checkbox"/> <code>hadoop_hdfs@%*</code>	Yes	GRANT USAGE ON *.* GRANT INSERT, UPDATE, DELETE, CREATE, DROP		
<input checked="" type="checkbox"/> <code>hadoop_hdfs@%*</code>	Yes	GRANT USAGE ON *.* GRANT ASSEMBLY, DELETE, CREATE, DROP		
<input type="checkbox"/> <code>hadoop_hdfs@%*</code>	No	GRANT INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, EXECUTION		View Delete
<input type="checkbox"/> <code>hadoop_hdfs@%*</code>	No	GRANT INSERT, UPDATE, DELETE, CREATE, DROP, RELOAD, EXECUTION		View Delete
<input checked="" type="checkbox"/> <code>hadoop_hdfs@%*</code>	Yes	GRANT USAGE ON *.* GRANT SELECT ON hdfs.*		

Below the table, there are sections for 'Migrate Password' and 'Create Instance Along With Task Structure', both with 'Yes' selected. At the bottom, there are 'Next' and 'Cancel' buttons.

Table 4-9 Migration types and objects

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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-11 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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.

Parameter	Description
	<p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Migrate Incremental Accounts and Permissions	<p>Indicates whether to migrate incremental accounts and permissions during database migration.</p> <ul style="list-style-type: none">● Yes All incremental accounts and permissions will be migrated. The migration of incremental accounts and permissions may fail because the source and destination database versions and account encryption modes may be different.● No All incremental accounts and permissions will be filtered out during the migration.
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none">● Yes If you need to migrate accounts, see Migrating Accounts.● No During migration, accounts, permissions, and passwords are not migrated.
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full migration phase.</p> <ul style="list-style-type: none">● Yes: Indexes are migrated when the table structure is migrated in the full migration phase.● No: Indexes are migrated separately after data migration.
Filter DROP DATABASE	<p>To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations.</p> <ul style="list-style-type: none">● If you select Yes, any database deletion operations performed on the source database are not migrated during data migration.● If you select No, related operations are migrated to the destination database during data migration.

Parameter	Description
Migration Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> • All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. • Tables: The selected table-level objects will be migrated. • Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. • 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 migration object cannot contain spaces. • To quickly select the desired database objects, you can use the search function.

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

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**

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

Step 6 Compare source and destination parameters.

By comparing common and performance parameters for the source databases against those of the destination databases, you can help ensure that services will not change after a migration is completed. You can determine whether to use this function based on service requirements. It mainly ensures that services are not affected after a migration is completed.

- This process is optional, so you can click **Next** to skip the comparison.
- Compare common parameters:

If the common parameter values in the comparison results are inconsistent, click **Save Change** to change the destination database values to be the same as those of the source database.

Figure 4-12 Modifying common parameters

Parameter Name	Source Database Value	Destination Database Value	Result
connect_timeout	10	10	Consistent
explicit_defaults_for_timestamp	OFF	OFF	Consistent
innodb_flush_log_at_trx_commit	1	1	Consistent
innodb_lock_wait_timeout	50	50	Consistent
innodb_buffer_pool_size	2500	153	Inconsistent
innodb_read_io_threads	30	30	Consistent
innodb_write_io_threads	60	60	Consistent
transaction isolation	REPEATABLE-READ	REPEATABLE-READ	Consistent

Performance parameter values in both the source and destination databases can be the same or different.

- If you need to change the performance parameter values that are consistent in the comparison results to different values, locate the target parameter, enter values in the **Change To** column, and click **Save Change** in the upper left corner.
- If you want to make the performance parameter values of the source and destination database be the same:
 - Click **Use Source Database Value**.
DRS automatically makes the destination database values the same as those of the source database.

Figure 4-13 One-click modification

Parameter Name	Source Database Value	Destination Database Value	Change To	Allowed Destination Database Value	Result	
innodb_buffer_pool_size	32768	32768	8	4096 - 32768	4096 - 10737216	Consistent
innodb_read_io_threads	32768	32768	8	4096 - 32768	4096 - 10737216	Consistent
innodb_write_io_threads	630800	630800		0 - 10487440733051615	0 - 10487440733051615	Consistent
innodb_buffer_pool_instances	2	2		1 - 64	1 - 64	Consistent
innodb_buffer_pool_size	429487296	429487296	16	256028096 - 429487296	1073741024 - 4817847873	Consistent
innodb_buffer_pool_instances	1300000	1300000	8	810 - 3008	810 - 3008	Consistent
innodb_read_io_threads	202144	202144	16	4096 - 202144	8192 - 214791902	Consistent
innodb_write_io_threads	524208	524208	128	4096 - 524208	4096 - 214791902	Consistent
innodb_buffer_pool_instances	202144	202144			32768 - 10487440733051615	Consistent
innodb_buffer_pool_instances	1	1			0 - 429487296	Consistent

NOTE

You can also manually enter parameter values.

- Click **Save Change** to save your changes.

The system changes the parameter values based on your settings for the destination database values. After the modification, the list is updated automatically.

Figure 4-14 Performance parameters

Parameter Name	Source Database Value	Destination Database Value	Change To	Allowed Destination Database V. (s)	Result
innodb_log_file_size	32768	32768	0	4096 + 32768	Consistent
innodb_log_files_in_group	32768	32768	0	4096 + 32768	Consistent
innodb_buffer_pool_size	8388608	8388608		0-16448144837370551615	Consistent
innodb_buffer_pool_instances	2	2		1-64	Consistent
innodb_buffer_pool_size	429497296	429497296	0	268435456 + 2147483648	Consistent
innodb_log_file_size	1300000	1300000		0-51-3000	Consistent
innodb_buffer_pool_size	262144	262144	0	4096 + 262144	Consistent
innodb_buffer_pool_size	524288	524288	0	4096 + 524288	Consistent
innodb_buffer_pool_size	262144	262144		32768-16448144837370551615	Consistent
innodb_log_file_size	1	1		0-429497296	Consistent

Some parameters in the destination database require a restart before the changes can take effect. The system will display these as being inconsistent. In addition, restart the destination database before the migration task is started or after the migration task is completed. To minimize the impact of this restart on your services, it is recommended that you schedule a specific time to restart the destination database after the migration is complete.

For details about how to set parameters during a comparison, see [Parameters for Comparison](#).

iii. Click **Next**.

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

Figure 4-15 Task startup settings

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

Apr 11, 2024 03:00:00 ⓘ GMT+08:00

In the full migration phase, the DRS task occupies performance resources. To ensure services and tasks can run normally, you are advised to start the task during off-peak hours and reserve two to three days for data verification.

Send Notifications: ⓘ

* SMN Topic: ⓘ ⓘ

Delay Threshold (s): ⓘ


Data Exception Notification:

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

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

Step 8 After the task is submitted, view and [manage it](#) on the **Online Migration 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.
- After the full migration is complete, you can use [data comparison](#) to check whether the data is consistent before and after the migration.
- 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.
- For a public network task, DRS needs to delete background resources after you stop the task. The EIP bound to the task cannot be restored to the **Unbound** state until background resources are deleted.

----End

Helpful Links

- [Supported Databases](#)
- [Preparations](#)
- [Migration Overview](#)

4.2 Querying the Migration Progress

The migration progress of a real-time migration task helps you keep track of the status of the migration task.

DRS shows the migration progress using a progress bar, helping you learn the migration progress in real time. During full migration, you can check migration details.

- In **Progress Overview**, you can view the migration progress of structures, data, and indexes. The migration progress is displayed based on the number of migrated objects. When the progress reaches 100%, the migration is complete. The migration of data and indexes is relatively slow during the migration.
- In the migration details, you can view the migration progress of a specific object. If the number of objects is the same as that of migrated objects, the migration is complete. You can view the migration progress of each object in detail. During incremental migration, the progress details are not displayed. You can view the consistency status on the **Migration Comparison** tab.

NOTE

- You can view migration details of all real-time migration tasks except Redis and GeminiDB Redis tasks.
- Before the task is completed, do not modify the information about all source and destination database users, passwords, and permissions.
- If the status indicates a synchronization or migration is complete, there may still be triggers or events to be migrated before the entire task is finished.

Prerequisites

- You have logged in to the DRS console.
- A migration task has been started.

Procedure

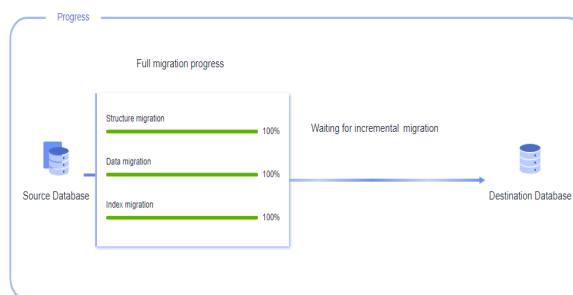
Step 1 On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 2 On the displayed page, click **Migration Progress**.

- View the migration progress of structures, data, and indexes.
When a full migration is complete, the progress of each item reaches 100%.
For a full plus incremental migration, you can view the delay of the incremental migration on the **Migration Progress** page.

Figure 4-16 Migration progress overview

Start the task and get 7 days free. If the status indicates a synchronization or migration is complete, there may still be triggers or events to be migrated before the entire task is finished.
Last Updated May 30, 2024 16:05:54 GMT+08:00 Source Position:mysql-0m-03365813675 Consumed Position N/A



You can also view the incremental migration delay on the **Online Migration Management** page. When the incremental migration delay exceeds the preset or default threshold, the value of the incremental migration delay is displayed in red in the task list.

NOTE

"Delay" refers to the delay from when the transaction was submitted to the source database to when it is synchronized to the destination database and executed.

Transactions are synchronized as follows:

1. Data is extracted from the source database.
2. The data is transmitted over the network.
3. DRS parses the source logs.
4. The transaction is executed on the destination database.

If the delay is 0, the source database is consistent with the destination database, and no new transactions need to be synchronized.

CAUTION

Frequent DDL operations, ultra-large transactions, and network problems may result in excessive synchronization delay.

- View the migration task progress. In the **Migration Details** area, locate the target migration object and click **View Details** in the **Operation** column to view the migration progress.

 **NOTE**

For a migration task from MySQL, the migration details are displayed during incremental migration. The details only display the data migrated in the full migration and the skipped structures and cannot be updated. For other migration tasks, the migration details are not displayed after incremental migration starts. You can compare migration items on the **Migration Comparison** page.

Figure 4-17 Migration progress details

Migration Details

During incremental migration, you can view the migration details on the Migration Comparison page.

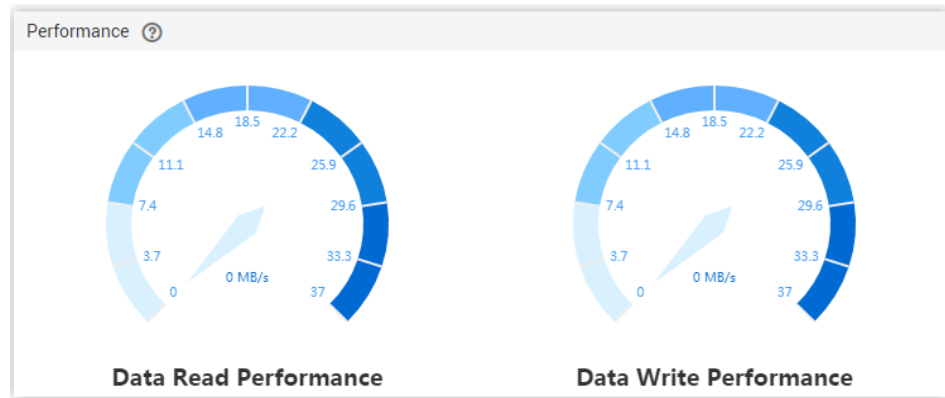
Migration Object	Total Items	Status	Migrated Items	Operation
account	5	Completed	5	View Details
database	1	Completed	1	View Details
event	0	Completed	0	View Details
function	0	Completed	0	View Details
procedure	0	Completed	0	View Details
table	2	Completed	2	View Details
table_index	0	Completed	0	View Details
table_structure	2	Completed	2	View Details
view	0	Completed	0	View Details
trigger	1	Pending migration	0	View Details

- Skip migration objects. In the **Migration Details** area, locate the migration object and click **View Details** in the **Operation** column. On the displayed page, click **Skip and Restart** in the **Operation** column to skip the object.

 **NOTE**

- For a migration task from MySQL, migration objects can be skipped during the full migration.
 - Migration objects that can be skipped include databases, functions, procedures, table_indexes, table_structure, views, triggers, events and accounts.
 - After a migration object is skipped, the task will be resumed. If the task is paused or abnormal, the task will be restarted after you click **Skip and Restart**.
 - If a task is in the full migration state, an object may be migrated before being skipped. In this case, skipping the object is not applied.
 - After an object is skipped, if operations associated with the object are still performed during subsequent migration, the task fails. For example, after table_structure is skipped, the task fails because the object is not found when indexes and data of the table are migrated.
 - After an object is skipped, the result for a comparison is inconsistent.
 - Skipping table_structure of a table without a primary key will cause data inconsistency. In addition, the task may fail because the table is not found. Perform the operation after confirmation.
 - If a DRS task fails because the table is not found after you skip a table without a primary key, contact O&M engineers for the destination database to create the corresponding table in the destination database based on the table structure of the source database and error log information, and click **Resume** in the **Operation** column of the task to submit the task again.
- Check the data read and write performance. Click the **Monitoring Graphs** to view the read and write performance. This graph shows the real-time rates of reading data from the source database and writing data to the destination database. The unit is MB/s.

Figure 4-18 Performance monitoring



----End

4.3 Viewing Migration Logs

Migration logs refer to the warning-, error-, and info-level logs generated during the migration process. This section describes how to view migration logs to locate and analyze database problems.

Prerequisites

- You have logged in to the DRS console.
- A migration task has been created.

Procedure

- Step 1** On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the **Migration Logs** tab, view logs of the migration task by level.

Figure 4-19 Viewing migration logs

Time	Level	Description
2024/04/16 17:59:48 GMT+08:00	info	increment transfer start
2024/04/16 17:59:48 GMT+08:00	info	set total migration stop point: mysql bin 025085191-c9231ec-ac35-11ea-af60-ba1634206d1-1-12867225
2024/04/16 17:59:48 GMT+08:00	info	full transfer completed
2024/04/16 17:59:48 GMT+08:00	info	flush privileges
2024/04/16 17:59:22 GMT+08:00	info	full transfer start
2024/04/16 17:59:22 GMT+08:00	info	set increment start point: mysql bin 025085191-c9231ec-ac35-11ea-af60-ba1634206d1-1-12867225
2024/04/16 17:52:58 GMT+08:00	info	precheck [a6c933c-eb60-48ab-887c-6c8a4a9191] completed, total item 41, success item 41, not pass item 1
2024/04/16 17:52:35 GMT+08:00	warning	The source mysql db have events or triggers.
2024/04/16 17:52:35 GMT+08:00	info	source charset is: utf8, target charset is: utf8
2024/04/16 17:52:24 GMT+08:00	info	precheck [a6c933c-eb60-48ab-887c-6c8a4a9191] start

You can view time, levels, and descriptions of the logs.

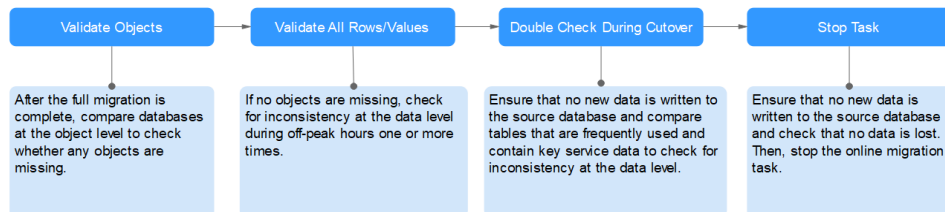
In addition, DRS can interconnect with Log Tank Service (LTS). After you enable log reporting to LTS, all logs generated by DRS instances will be uploaded to LTS for management. For details, see [Log Reporting](#).

----End

4.4 Data Comparison (Comparing Migration Items)

This section describes how to compare migration items to check if there are any differences between source and destination databases. By comparing migration objects, you can determine the proper time for service migration to minimize the service downtime.

Figure 4-20 Migration comparison process



Comparison Scenarios

You can compare migration objects with different dimensions:

- Object-level comparison: It helps you compare databases, indexes, tables, views, stored procedures and functions, and sorting rules of tables. You are advised to perform the comparison after a full migration is complete.
- 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 migrated. This comparison method is recommended because it is fast.
 - Value comparison: It helps you check whether data in the migrated table is consistent. The comparison process is relatively slow.
 - Sampling comparison: If there is a lot of data, it takes a long time to compare the number of rows and values. Sampling comparison is recommended. You can set an appropriate sampling ratio to complete a comparison faster.
- Account comparison: It compares usernames and permissions of the source and destination databases.
- Periodic comparison: DRS periodically compares the number of rows or objects in the source database table with those in the destination database table and displays the comparison results. To compare rows or 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 4-11 Supported comparison methods

Migratio n 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 cloud	MySQL->MySQL	Supported	Supported	Supported	Supported	Supported
To the cloud	MySQL->GaussDB(for MySQL)	Supported	Supported	Supported	Supported	Supported
To the cloud	MySQL->DDM	Supported	Supported	Not supported	Not supported	Not supported
To the cloud	MongoDB->DDS	Supported	Supported	Supported	Not supported	Supported
To the cloud	MongoDB->GeminiDB Mongo	Supported	Supported	Supported	Not supported	Not supported
To the cloud	MySQL schema and logic table -> DDM	Supported	Supported	Not supported	Not supported	Not supported
To the cloud	Redis->GeminiDB Redis	Not supported	Not supported	Not supported	Not supported	Not supported
To the cloud	Redis Cluster -> GeminiDB Redis	Not supported	Not supported	Not supported	Not supported	Not supported
From the cloud	MySQL->MySQL	Supported	Supported	Supported	Supported	Supported

Migration Direction	Data Flow	Object-level Comparison	Row Comparison	Value Comparison	Dynamic Comparison	Account-level Comparison
From the cloud	DDS->MongoDB	Supported	Supported	Supported	Not supported	Not supported
From the cloud	GeminiDB Redis->Redis	Not supported	Not supported	Not supported	Not supported	Not supported
From the cloud	GeminiDB Redis->Redis	Not supported	Not supported	Not supported	Not supported	Not supported

Comparison Restrictions

- You can manually create a comparison task only when the task is in the incremental phase.
- During a comparison, the comparison items are case sensitive. If one of the source or destination database is case insensitive and the other one is case sensitive, the comparison result may be inconsistent.
- When a full migration task is complete, DRS automatically creates object-level and row comparison tasks. If operations are performed on data in the source database during data comparison, the comparison results may be inconsistent.
- If DDL operations were performed on the source database, you need to compare the objects again to ensure the accuracy of the comparison results.
- If data in the destination database is modified separately, the comparison results may be inconsistent.
- If the encoding of the source database character type is abnormal, the database driver will convert the character type to an abnormal code point during DRS migration or comparison. As a result, the values may be consistent but the bytes may be inconsistent.
- Currently, only tables with primary keys support value comparison. For tables that do not support value comparison, you can compare rows. Therefore, you can compare data by row or value based on scenarios.
- The DRS task cannot be suspended during value comparison. Otherwise, the comparison task may fail.
- Some data types do not support value comparison. For details, see [Which of the Following Data Types Are Not Supported By Value Comparison?](#)

- To prevent resources from being occupied for a long time, DRS limits the comparison duration. If the comparison duration exceeds the threshold, the comparison task stops automatically.
 - When a full migration task is complete, DRS automatically creates object-level and row comparison tasks. The comparison duration limits to 30 minutes. After the threshold, the comparison tasks automatically stop and the full migration task stops.
 - For a row comparison task manually created in the incremental phase, if the source database is a relational database, the row comparison duration limits to 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration limits to 30 minutes.
- To avoid occupying resources, the comparison results of DRS tasks can be retained for a maximum of 60 days. After 60 days, the comparison results are automatically cleared.
- For a migration task from MySQL, virtual columns in the source database do not support value comparison. During the comparison, virtual columns are filtered out.

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 migration task has been started.

Creating a comparison task

You can follow the comparison process or select a comparison method based on your service scenario. The following operations describe how to compare migration items by following the recommended migration process.

Step 1 On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 2 On the **Migration Comparison** tab, compare objects of the source and destination databases.

You can also select the migration task on the **Online Migration Management** page and click **View** to go to the **Migration Comparison** page.

- Create an **Object-Level Comparison** task to check the integrity of database objects.


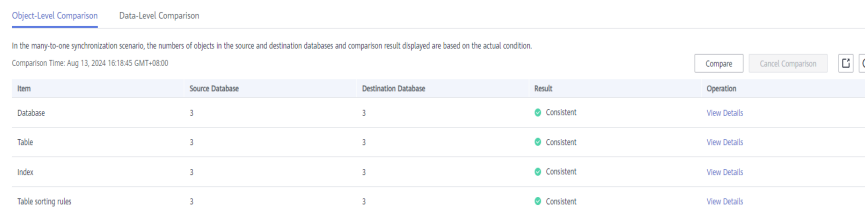
Click **Object-Level Comparison**. On the **Object-Level Comparison** tab, click **Compare**. Wait for a while and click , and view the comparison result of each comparison item.

Figure 4-21 Object-Level Comparison



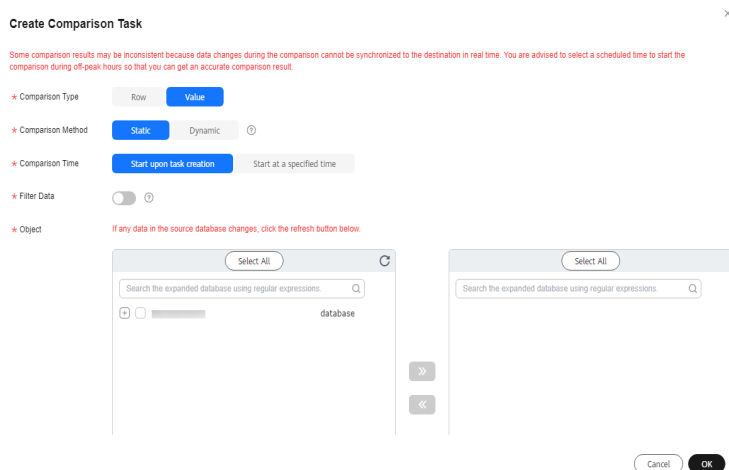
Item	Source Database	Destination Database	Result	Operation
Database	3	3	Consistent	View Details
Table	3	3	Consistent	View Details
Index	3	3	Consistent	View Details
Table sorting rules	3	3	Consistent	View Details

Locate a comparison item you want to view and click **View Details** in the **Operation** column.

- Create a **Data-Level Comparison (row comparison or value comparison)** task to compare the number of rows and values of the migrated data.

If you only need to compare the number of rows of all migration objects, you can select a specified migration task on the **Online Migration Management** page and click **Compare** in the **Operation** column to create a comparison task.

- In the **Before You Start** pane, click **Validate All Rows/Values**.
- In the displayed **Create Comparison Task** dialog box, specify **Comparison Type**, **Comparison Method**, **Comparison Time**, and **Object**. Then, click **OK**.

Figure 4-22 Creating a comparison task

- **Comparison Type:** compares rows and values.
- **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

- The comparison mode can only be changed for MySQL.
- During database-level migration, 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 migrated, 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.
- **Filter Data:** After this function is enabled, objects can be compared based on the configured filtering criteria.

NOTE

Only MySQL-to-MySQL migration tasks support data filtering and comparison.

After enabling **Filter Data**, add filtering criteria for the table objects to be compared.

In the **Filtering Criteria** area, enter the filtering criteria, and click **Verify**.

 **NOTE**

- Standard SQL statements can be used to filter records. Each expression cannot contain packages, functions, variables, or constants specific to a database engine.
- Enter the part following WHERE in the SQL statement (excluding WHERE and semicolons), for example, sid > 3 and sname like "G %".
- 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.
- You are not advised to set filter criteria for fields of approximate numeric types, such as FLOAT, DECIMAL, and DOUBLE.
- Do not use fields containing special characters as a filter condition.
- Objects whose database names, schema names, or table names are case insensitive cannot be filtered and compared.
- Currently, condition-based filtering is not supported when there are more than 50,000 tables in a database.


After the verification is successful, click **Generate Processing Rule**. The rule is displayed.

Click **OK**.

- **Object:** You can select objects to be compared based on the scenarios.

 **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.

- c. 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.

To view the comparison details, locate the target comparison type and click **View Results** in the **Operation** column. On the displayed page, locate a pair of source and destination databases, and click **View Details** in the **Operation** column to view detailed comparison results.

 NOTE

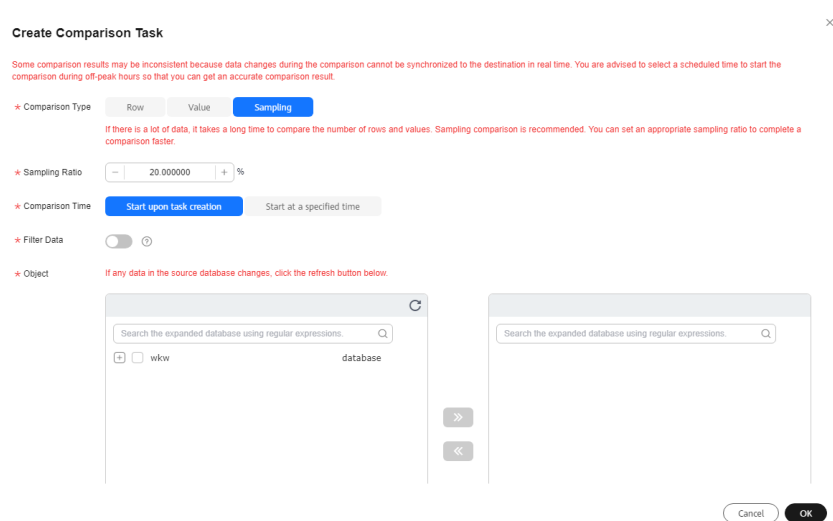
- You can cancel a running task at any time and view the comparison report of a canceled comparison task.
 - You can sort the row comparison results displayed on the current page in ascending or descending order based on the number of rows in the source database table or the destination database table.
 - If a negative number is displayed in the **Row 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.
- Create a **Data-Level Comparison (sampling comparison)** task. If there is a lot of data, it takes a long time to compare the number of rows and values. Sampling comparison is recommended. You can set an appropriate sampling ratio to complete a comparison faster.

On the **Data-Level Comparison** page, click **Create Comparison Task**, select **Sample** for **Comparison Type**, set **Sampling Ratio**, **Comparison Time**, and **Object**, and click **OK**.

 NOTE

- Only real-time migration tasks from MongoDB to DDS and from DDS to MongoDB support sampling comparison.
- The **\$sample** command provided by MongoDB is used for data sampling: **{ \$sample: { size: <positive integer N> } }**.
In the preceding command, **size** is fixed to 1000. To perform MongoDB sampling comparison provided by DRS, the **\$sample** command is executed for multiple times until the sampled data volume reaches the specified sampling ratio.
- In a single sampling comparison task, the **\$sample** command is executed for multiple times. In this case, the same data in the source database may be obtained for multiple times.
- The total data volume of a collection is obtained using **estimatedDocumentCount**. It is an estimated value instead of an accurate value. The actual sampling data volume and sampling ratio may fluctuate, which is normal.
- Sampling comparison does not support documents whose **_id** is of the BinData data type. Documents whose **_id** is of the BinData data type will be filtered out and not be compared.
- If the number of inconsistent data records in a single MongoDB sampling comparison task exceeds 10,000, the comparison stops. A maximum of 10,000 inconsistent data records can be displayed.

Figure 4-23 Sampling comparison



- Create an **Account-Level Comparison** task to compare database accounts and permissions. Click the **Account-Level Comparison** tab to view the comparison results of database accounts and permissions.

Figure 4-24 Account-level comparison

Source Database Account Attribute	Source Database Account Name	Destination Database Account Attribute	Destination Database Account Name	Migration Comparison Time	Result
CREATOR,CREATOR,LEN,NONHERITPASSW...	pkuser1	CREATOR,CREATOR,LEN,NONHERITPASSW...	pkuser1	2021/06/19 11:39:28 GMT+08:00	Consistent

NOTE

- Full migration tasks do not support account-level comparisons.
- Perform a double check before the cutover. Click **Double Check During Cutover**. In the displayed **Create Comparison Task** dialog box, specify **Comparison Type**, **Comparison Time**, and **Object**. Then, click **OK**. For details about how to view comparison details, see [Step 2.c](#).
- Stop the migration task.

After the service system is successfully migrated to the destination database, stop the migration task to prevent operations in the source database from being synchronized to the destination database to overwrite the data. This operation only deletes the replication instance, and the migration task is still in the task list. You can view or delete the task. DRS will not charge for this task after you stop it.

Generally, stopping a task can ensure the integrity of special objects because triggers and events are migrated when a task is being stopped. Only in some cases, such as network disconnections, a task may fail to be stopped. If a task fails to be stopped multiple times, you can select **Forcibly stop task** to reduce the waiting time. If you forcibly stop a task, triggers and events may not be completely migrated and you need to manually migrate them.

----End

Periodic Comparison

Periodic comparison indicates that DRS periodically compares the number of rows or objects in the source database table with those in the destination database table and displays the comparison results.

Step 1 On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 2 Choose **Migration Comparison**.

Step 3 Click the **Periodic Comparison** tab and click **Modify Comparison Policy**.

NOTE

Only MySQL-to-MySQL and MySQL-to-GaussDB(for MySQL) migration tasks support periodic comparison.

Step 4 In the **Modify Comparison Policy** dialog box, enable periodic comparison, specify the comparison frequency, time interval, effective time and comparison type, and click **Yes**.

NOTE

- After periodic comparison is enabled, DRS compares the number of rows or objects at the scheduled time. You can view the comparison results on the **Data-Level Comparison** or **Object-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 4-25 Modify Comparison Policy

----End

Quick Comparison

To accelerate and simplify the migration process, DRS provides the quick comparison function. You can directly perform a comparison on the migration task list. This function can be used to compare all migration objects only when incremental migration tasks are in progress.

- Step 1** On the **Online Migration Management** page, locate the target migration task and click **Compare** in the **Operation** column.
- Step 2** On the **Create Comparison Task** page, select **Start upon task creation** or **Start at a specified time** and click **Yes** to start the comparison task.

----End

Viewing a Comparison Task

- Step 1** On the **Online Migration Management** page, locate the target migration task and click **View** in the **Operation** column.

Step 2 On the **Migration Comparison** tab, view the data comparison result.

----End

4.5 Managing Objects

4.5.1 Migrating Accounts

Scenarios

During a database migration, accounts need to be migrated separately.

MySQL Databases Operations

During the migration of MySQL databases, there are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated.

- Accounts that can be completely migrated refer to the accounts that meet the permission requirements of the destination database. By default, the system automatically migrates the permission of the database account to the destination database.
- Accounts whose permissions need to be reduced refer to high-level accounts that fail to meet the permission requirements of the destination database, such as super, file, and shutdown. To migrate these accounts, reduce the permissions of the account. Otherwise, the migration fails.

You can click **View** in the **Remarks** column to view detailed information about the permission to be reduced. You can then determine whether the permission reduction will have an impact on your services.

- Accounts that cannot be migrated indicate that database users cannot meet the migration requirements for certain reasons. These accounts will not be migrated to the destination database. Ensure that services are not affected by these accounts. After the migration is started, any operation of changing the password or permission for these accounts will result in an incremental migration failure.

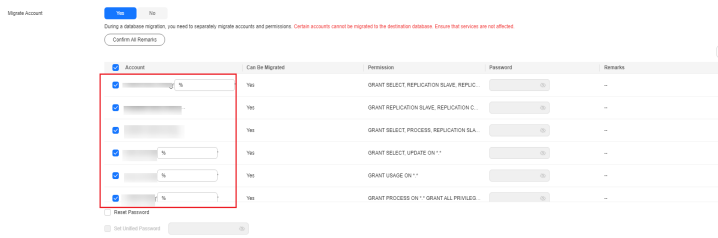
You can choose whether to migrate the accounts or choose to migrate certain of or all of the accounts. Perform the following operations to set the database username, permission, and password. The following procedure uses all database users that can be migrated as an example.

The account information consists of account name, permission, and password.

Step 1 The account name is in the '**Account name**'+'@+'**host**' format. **host** indicates the IP address of the destination database, which is allowed to access the source database. You can change the IP address as required.

The IP address in the red box is shown in the following figure. If the source database is MySQL 8.0, you are not advised to change the IP address.


Figure 4-26 Changing the IP address



Step 2 By default, account permissions cannot be modified. For accounts that can be migrated (including accounts that can be completely migrated and accounts whose permissions need to be reduced), the system also migrates the permissions of these accounts.

After the migration is successful, accounts in the destination database are those whose permissions need to be reduced.

Figure 4-27 Account permissions



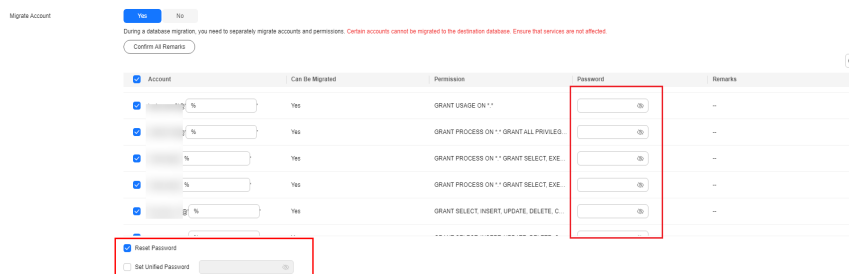
Step 3 Migrate account passwords.

You can use either of the following methods to migrate account passwords:

DRS does not check your password strength during migration so you should set a strong password to ensure data security.

Method 1: Migrate the password.

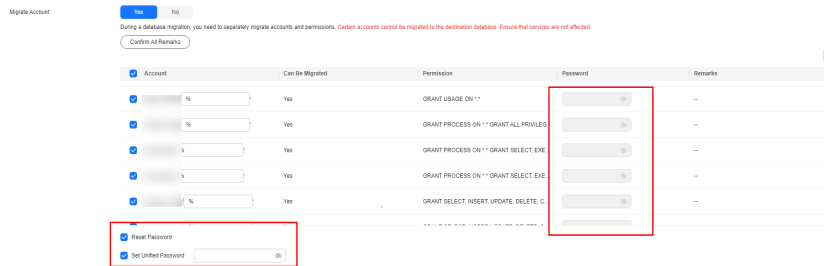
Figure 4-28 Migrating password



You can directly migrate the current password of the source system. In this case, you do not need to select **Reset Password**. After the passwords are migrated to the destination database, you can set a strong password to ensure database security.

Method 2: Reset the password.

Figure 4-29 Resetting a password

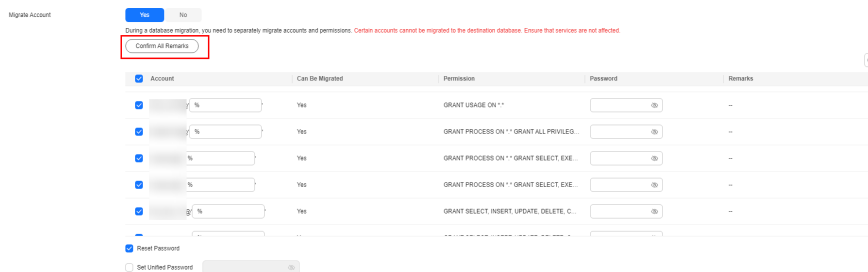


You can select **Reset Password** to reset the password of the source system and then continue the password migration.

You can enter new passwords in the **Passwords** column for specified accounts that can be migrated, or select all accounts that can be migrated and select **Set Unified Password** to set a unified new password for them. After the migration is successful, you can run DDL statements on the destination database to reset the password.

- Step 4** For accounts whose permissions need to be reduced and accounts that cannot be migrated, you can click **View** to confirm the remarks before performing the next step. If there are multiple accounts, you can click **Confirm All Remarks**.

Figure 4-30 Remarks



If an account already exists in the destination database, it cannot be migrated. You can delete it from the destination database. After the deletion, you can continue the migration.

NOTE

- The new password you set must meet the **password policy of the destination database**.

----End

MongoDB Database Operations

During the migration of MongoDB databases, accounts to be migrated can be classified into the following types: accounts that can be migrated completely and accounts that cannot be migrated.

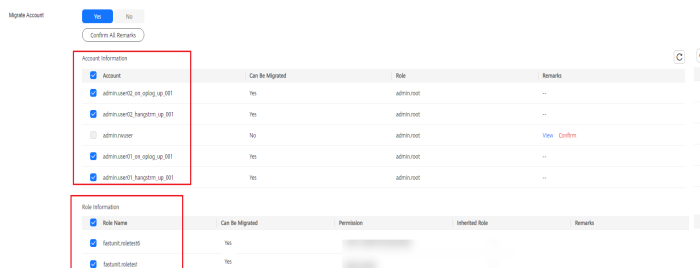
You can choose whether to migrate the accounts. If you need to migrate the accounts, perform the following procedures.

The account information consists of account name and role.

Step 1 Select the accounts and roles to be migrated based on service requirements.


If the account to be migrated depends on some roles, you must migrate the roles. Otherwise, the migration fails.

Figure 4-31 Account migration



Step 2 For accounts or roles that cannot be migrated, you can click **View** to confirm the remarks before performing the next step. If there are multiple accounts, you can click **Confirm All Remarks**.

Figure 4-32 Remarks



----End

4.5.2 Parameters for Comparison

Parameter comparison helps you check consistency between the source and destination database data to ensure your services will not be affected after being migrated.

This section lists the common parameters and performance parameters of different DB engine versions for your reference during parameter comparison.

MySQL 5.6

Table 4-12 MySQL 5.6 parameters to be compared

Parameter	Type	Restart Required
connect_timeout	Common parameter	No
event_scheduler	Common parameter	No
innodb_lock_wait_timeout	Common parameter	No
max_connections	Common parameter	No
net_read_timeout	Common parameter	No
net_write_timeout	Common parameter	No
explicit_defaults_for_timestamp	Common parameter	Yes
innodb_flush_log_at_trx_commit	Common parameter	No
max_allowed_packet	Common parameter	No
tx_isolation	Common parameter	No
character_set_client	Common parameter	No
character_set_connection	Common parameter	No
collation_connection	Common parameter	No
character_set_results	Common parameter	No
collation_server	Common parameter	No
binlog_stmt_cache_size	Performance parameter	No
bulk_insert_buffer_size	Performance parameter	No
innodb_buffer_pool_size	Performance parameter	Yes
innodb_buffer_pool_instances	Performance parameter	Yes
key_buffer_size	Performance parameter	No
long_query_time	Performance parameter	No
query_cache_type	Performance parameter	Yes
read_buffer_size	Performance parameter	No
read_rnd_buffer_size	Performance parameter	No
sort_buffer_size	Performance parameter	No

Parameter	Type	Restart Required
sync_binlog	Performance parameter	No

MySQL 5.7

Table 4-13 MySQL 5.7 parameters to be compared

Parameter	Type	Restart Required
connect_timeout	Common parameter	No
event_scheduler	Common parameter	No
innodb_lock_wait_timeout	Common parameter	No
max_connections	Common parameter	No
net_read_timeout	Common parameter	No
net_write_timeout	Common parameter	No
explicit_defaults_for_timestamp	Common parameter	No
innodb_flush_log_at_trx_commit	Common parameter	No
max_allowed_packet	Common parameter	No
tx_isolation	Common parameter	No
character_set_client	Common parameter	No
character_set_connection	Common parameter	No
collation_connection	Common parameter	No
character_set_results	Common parameter	No
binlog_cache_size	Performance parameter	No
binlog_stmt_cache_size	Performance parameter	No
bulk_insert_buffer_size	Performance parameter	No
innodb_buffer_pool_size	Performance parameter	No
innodb_buffer_pool_instances	Performance parameter	Yes
key_buffer_size	Performance parameter	No
long_query_time	Performance parameter	No
query_cache_type	Performance parameter	No

Parameter	Type	Restart Required
read_buffer_size	Performance parameter	No
read_rnd_buffer_size	Performance parameter	No
sort_buffer_size	Performance parameter	No
sync_binlog	Performance parameter	No

MySQL 8.0

Table 4-14 MySQL 8.0 parameters to be compared

Parameter	Type	Restart Required
connect_timeout	Common parameter	No
event_scheduler	Common parameter	No
innodb_lock_wait_timeout	Common parameter	No
max_connections	Common parameter	No
net_read_timeout	Common parameter	No
net_write_timeout	Common parameter	No
explicit_defaults_for_timestamp	Common parameter	No
innodb_flush_log_at_trx_commit	Common parameter	No
max_allowed_packet	Common parameter	No
tx_isolation	Common parameter	No
character_set_client	Common parameter	No
character_set_connection	Common parameter	No
collation_connection	Common parameter	No
character_set_results	Common parameter	No
binlog_cache_size	Performance parameter	No
binlog_stmt_cache_size	Performance parameter	No
bulk_insert_buffer_size	Performance parameter	No
innodb_buffer_pool_size	Performance parameter	No
innodb_buffer_pool_instances	Performance parameter	Yes

Parameter	Type	Restart Required
key_buffer_size	Performance parameter	No
long_query_time	Performance parameter	No
query_cache_type	Performance parameter	No
read_buffer_size	Performance parameter	No
read_rnd_buffer_size	Performance parameter	No
sort_buffer_size	Performance parameter	No
sync_binlog	Performance parameter	No

NOTE

- Currently, only MySQL databases support the parameter comparison function.
- The value of **innodb_buffer_pool_size** is set to not exceed 70% of the total memory of the destination database. If you set a larger value for parameter, the destination database startup may fail. Therefore, values of **innodb_buffer_pool_size** in the source and destination databases are different. You can adjust the value to suit your services.

4.5.3 Viewing Migration Objects

After a migration task is configured and started, you can view the migration object details.


Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the displayed page, choose **Object Details** to view the migration objects of the migration task.

Figure 4-33 Viewing migration objects



 **NOTE**

When an object is displayed, 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.

----End

4.6 Task Life Cycle

4.6.1 Viewing Task Details

This section describes how to view details about a migration task, including information about the task, replication instance, and migration.

Prerequisites

- You have logged in to the DRS console.
- A migration task has been created.

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 **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 2 On the displayed **Basic Information** tab, view details about the migration task.

You can view information about the task, replication instance, and migration.

----End

4.6.2 Modifying Task Information

After a migration 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.

- A migration task has been created.

Procedure

Step 1 On the **Online Migration Management** page, click the target migration 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 4-15 Task information

Task Information	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain the following special characters: !<>&'\"
SMN Topic	You can apply for a topic on the SMN console and add a subscription. For details, see Simple Message Notification User Guide .
Synchronization Delay Threshold	The delay ranges from 0s to 3600s. NOTE If the delay threshold is set to 0, no notifications will be sent to the recipient.
Stop Abnormal Tasks After	The value must range from 14 to 100. The default value is 14.

- You can modify the task start time only when the task is in the **Pending start** status.

In the **Task Information** area, click **Modify** in the **Scheduled Start Time** field. On the displayed page, specify the scheduled start time and click **OK**.

Step 3 View the change result on the **Basic Information** tab.

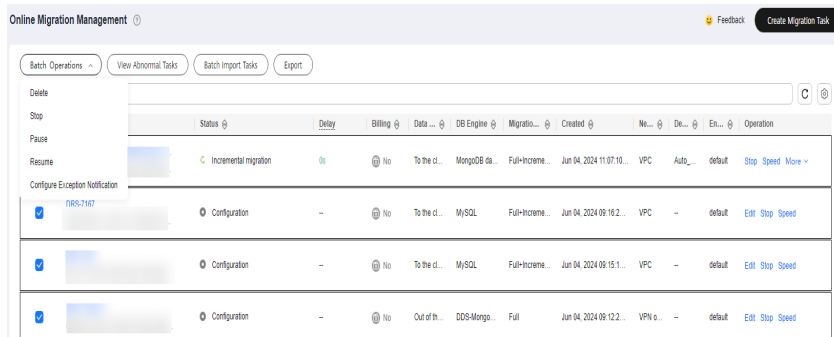
----End

Configuring Exception Notifications

Step 1 On the **Online Migration Management** page, select the task to be configured.

Step 2 Click **Batch Operations** in the upper left corner and choose **Configure Exception Notification**.

Figure 4-34 Batch Operations



Step 3 In the displayed dialog box, enter the configuration information and click **Yes** to submit the configuration task.

----End

4.6.3 Modifying Connection Information

During the migration, you may change the password of the source or destination database. As a result, the data migration, data comparison, task pause, resume, resetting, 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 information:

- Database password
- Database IP address
- Database port
- Database username
- DNS server

Constraints

- The database connection password can be changed for all DRS tasks.
- You can change the IP address, port, and username during the incremental migration phase only for a DRS task with 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 **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click **Modify Connection Details** in the **Migration Information** area.
- Step 3** In the displayed dialog box, change the database information such as passwords of the source and destination databases and click **OK**.
- Step 4** If a task is in the **Failed** state, return to the task list page after the change is complete and click **Resume** in the **Operation** column to continue the migration task.

----End

4.6.4 Modifying the Flow Control Mode

You can choose whether to control the flow. DRS allows you to change the flow control mode after a task is created. Currently, only the following real-time migration types support this function:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL->DDM
 - MySQL schema and logic table -> DDM
 - MongoDB->DDS
 - MongoDB->GeminiDB Mongo
- From of the cloud
 - MySQL->MySQL
 - DDS->MongoDB

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 migration 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.

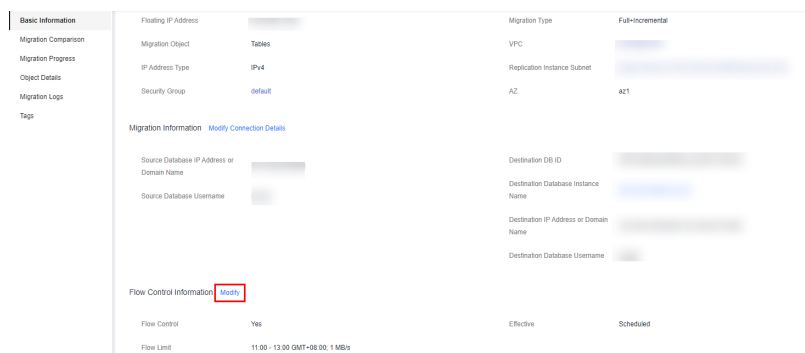
Prerequisites

- You have logged in to the DRS console.
- A migration task has been created.

Method 1

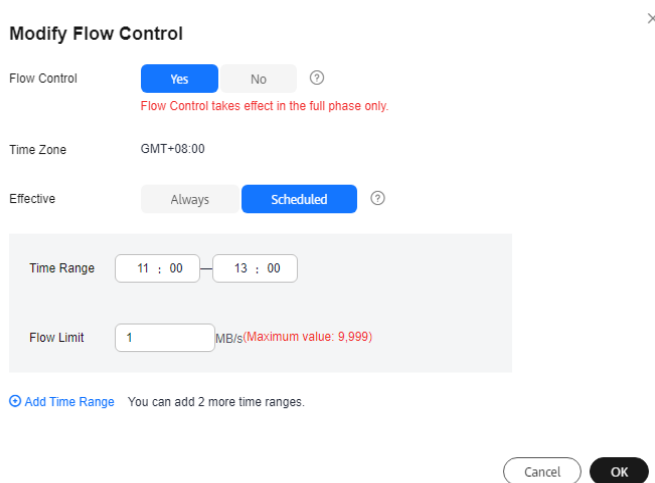
- Step 1** In the **Flow Control Information** area on the **Basic Information** tab, click **Modify**.

Figure 4-35 Basic information



Step 2 In the displayed dialog box, modify the settings.

Figure 4-36 Modify Flow Control




----End

Method 2

Step 1 In the task list on the **Online Migration Management** page, locate the target task and choose **More > Speed** or **Speed** in the **Operation** column.

Figure 4-37 Task list



Step 2 In the displayed dialog box, modify the settings.

Figure 4-38 Modify Flow Control

Modify Flow Control ×

Flow Control Yes No ⓘ
Flow Control takes effect in the full phase only.

Time Zone GMT+08:00

Effective Always Scheduled ⓘ

Time Range 11 : 00 — 13 : 00

Flow Limit 1 MB/s (Maximum value: 9,999)

[Add Time Range](#) You can add 2 more time ranges.

Cancel OK

----End

4.6.5 Editing a Migration Task

For a migration task that has been created but not started, DRS allows you to edit the configuration information of the task, including the task information, replication instance information, and migration information. For migration tasks in the following statuses, you can edit the tasks again after the replication instances are created:

- Creating
- Configuration

NOTE

For a started migration task, modifying the migration objects is not supported.

Prerequisites

- You have logged in to the DRS console.
- A migration task has been created.

Method 1

- Step 1** In the task list on the **Online Migration 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 Task** page, select the accounts and objects to be migrated, and click **Next**.

Figure 4-39 Migration type

Filter Criteria

Migrate Instance Accounts and Permissions The migration of Instance accounts and permissions may fail because the source and destination RDS instances and account migration modes may be different.

Migrate Account During cross-instance migration, you need to migrate user accounts and permissions. User accounts created in the destination database. Except for users that are not allowed. A list of user accounts can be required. After RDS accounts are migrated to the cloud, the migrated account is displayed in the table below. View Details

[Confirm All Results](#)

<input checked="" type="checkbox"/> Account	Can Be Migrated	Permissions	Password	Remarks
<input checked="" type="checkbox"/> sysadmin_001@% .%	Yes	GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* GRANT SELECT	<input type="button" value="No"/>	-
<input checked="" type="checkbox"/> sys_replicat_user@% .%	Yes	GRANT USAGE ON *.* GRANT REPLICATION SLAVE, DELETE, CREATE, DROP R.	<input type="button" value="No"/>	-
<input checked="" type="checkbox"/> sys_replicat_slave@% .%	Yes	GRANT USAGE ON *.* GRANT REPLICATION SLAVE, DELETE, CREATE, DROP R.	<input type="button" value="No"/>	-
<input type="checkbox"/> sys_admin_user@% .%	No	GRANT REPLICATION SLAVE, CREATE, CREATE, DROP, RELOAD, SHUTDOWN,	-	View Details
<input type="checkbox"/> sys_replicat_slave@% .%	No	GRANT REPLICATION SLAVE, CREATE, CREATE, DROP, RELOAD, SHUTDOWN,	-	View Details
<input checked="" type="checkbox"/> sys_replicat_slave@% .%	Yes	GRANT USAGE ON *.* GRANT SELECT ON *.* REPLICATION SLAVE	<input type="button" value="No"/>	-

Migrate Password

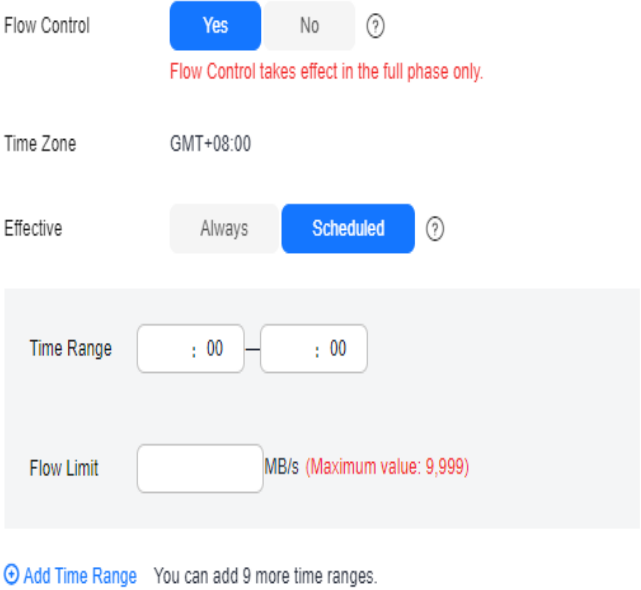
Set Global Password

Create Instance Using VMS Task Scheduler


Filter DMS DMS/DC/DC/DC

Migrate Object

Table 4-16 Migration types and objects

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 migration speed. During the migration, the migration speed of each task (or each subtask in multi-task mode) does not exceed the value of this parameter. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is Always. A maximum of 10 time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. <p>Figure 4-40 Configuring flow control</p>  <ul style="list-style-type: none"> No The migration 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.

Parameter	Description
	<p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Migrate Incremental Accounts and Permissions	<p>Indicates whether to migrate incremental accounts and permissions during database migration.</p> <ul style="list-style-type: none">● Yes All incremental accounts and permissions will be migrated. The migration of incremental accounts and permissions may fail because the source and destination database versions and account encryption modes may be different.● No All incremental accounts and permissions will be filtered out during the migration.
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none">● Yes If you need to migrate accounts, see Migrating Accounts.● No During migration, accounts, permissions, and passwords are not migrated.
Create Indexes Along With Table Structure	<p>Indicates whether to create indexes along with the table structure in the full migration phase.</p> <ul style="list-style-type: none">● Yes: Indexes are migrated when the table structure is migrated in the full migration phase.● No: Indexes are migrated separately after data migration.
Filter DROP DATABASE	<p>To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations.</p> <ul style="list-style-type: none">● If you select Yes, any database deletion operations performed on the source database are not migrated during data migration.● If you select No, related operations are migrated to the destination database during data migration.

Parameter	Description
Migration Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none"> ● All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified. ● Tables: The selected table-level objects will be migrated. ● Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none"> ● If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases. ● 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 migration object cannot contain spaces. ● To quickly select the desired database objects, you can use the search function.

Step 4 On the **Check Task** page, check the migration task.

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

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

- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**


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

Step 5 On the **Confirm Task** page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

 NOTE

- Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements.
- After a migration task is started, the performance of the source and destination databases may be affected. You are advised to start a migration task during off-peak hours.
- Under specific conditions, the destination database needs to be restarted once during the task startup, which may interrupt database services.

Step 6 After the task is submitted, view and manage it on the **Online Migration 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.

----End

Method 2

Step 1 On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 2 On the displayed page, click **edit this task** to go to the **Configure Source and Destination Databases** page.

Step 3 Perform steps [Step 2](#) to [Step 6](#).

----End

4.6.6 Resuming a Migration Task

A fault may occur during the migration due to external factors, such as insufficient storage space. After the fault is rectified based on the migration log information, you can resume the migration.

You can resume migration tasks in any of the following statuses:

- Migration failed
- Paused

You can resume the following migration tasks:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL->DDM
 - MySQL schema and logic table -> DDM
 - MongoDB->DDS
 - MongoDB->GeminiDB Mongo
- From the cloud
 - MySQL->MySQL
 - DDS->MongoDB

NOTE

- If a migration 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.

Prerequisites

- You have logged in to the DRS console.
- A migration task has been created.

Method 1

On the **Online Migration Management** page, locate the target task and click **Resume** in the **Operation** column.

Method 2

Step 1 In the task list on the **Online Migration Management** page, locate and click the task.

Step 2 On the displayed page, click the **Migration Progress** tab, and click **Resume** in the upper right corner.

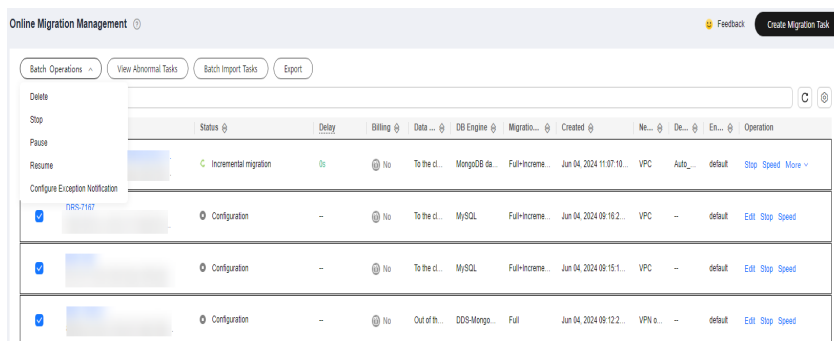
----End

Resuming Tasks

Step 1 On the **Online Migration Management** page, select the tasks to be resumed.

Step 2 Click **Batch Operations** in the upper left corner and choose **Resume**.

Figure 4-41 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

4.6.7 Disabling Read-Only of the Destination Database

For a running migration task, DRS allows you to disable read-only of the destination database.

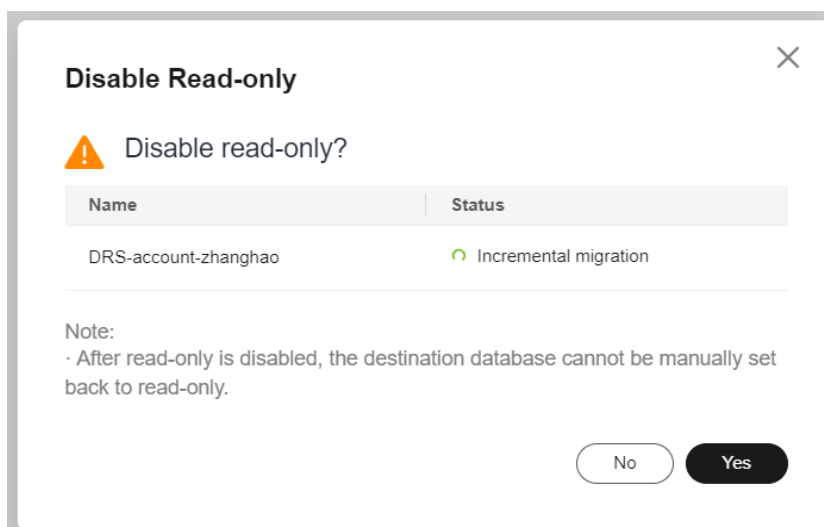
Constraints

- You can disable read-only of the destination database only for to-the-cloud migration tasks from MySQL to MySQL, MySQL to GaussDB(for MySQL), and MongoDB to DDS.
- Disabling read-only on the destination database may cause data inconsistencies. Exercise caution when performing this operation.
- After read-only is disabled, the destination database cannot be set back to read-only. Even if the migration task fails and then resumed and started, the destination database cannot be set to read-only.
- The access settings of the destination instance are also affected by the access settings of the DB instance itself. Therefore, you are advised not to set the access settings of the DB instance on the RDS console.

Disabling Read-Only

- Step 1** In the task list on the **Online Migration Management** page, locate the target task and click its name.
- Step 2** In the **Task Information** area on the **Basic Information** page, click **Disable Read-only** next to **Destination DB Instance Access**.
- Step 3** In the displayed dialog box, click **Yes**.

Figure 4-42 Disabling Read-Only



----End

4.6.8 Resetting a Migration Task

During the migration, if a migration task fails due to uncertain causes, the background will resume the task several times. However, the task may fail to be recovered in some scenarios. To continue the migration, DRS allows you to reset the task.

You can reset the following DRS tasks:

- MySQL->MySQL

- MySQL->DDM
- MySQL->GaussDB(for MySQL)
- MySQL schema and logic table -> DDM
- Redis->GeminiDB Redis
- Redis Cluster -> GeminiDB Redis
- MongoDB->DDS
- DDS->MongoDB
- GeminiDB Redis->Redis
- GeminiDB Redis -> Redis Cluster

Prerequisites

- You have logged in to the DRS console.
- A migration task has failed.

Method 1

Step 1 In the task list on the **Online Migration Management** page, locate the target task and click **Reset** in the **Operation** column.

Step 2 In the displayed dialog box, check the migration task again.

Step 3 After the check is complete and the check success rate is 100%, click **Start** to submit the migration task again.

----End

Method 2

Step 1 On the **Data Migration Management** page, click the target task name in the **Task Name/ID** column.

Step 2 On the displayed page, click the **Migration Progress** tab, and click **Reset** in the upper right corner.

Step 3 Perform [Step 2](#) to [Step 3](#) from method 1.

----End

4.6.9 Pausing a Migration Task

During migration, if the flow control mode cannot meet the requirements during peak hours, you can pause the migration task.

You can pause the following migration tasks:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(for MySQL)
 - MySQL->DDM
 - MySQL schema and logic table -> DDM

- MongoDB->DDS
- MongoDB->GeminiDB Mongo
- From the cloud
 - MySQL->MySQL
 - DDS->MongoDB

Prerequisites

- You have logged in to the DRS console.
- The migration task is running properly.

Pausing a Task

Step 1 In the task list on the **Online Migration Management** page, locate the target task and click **Pause** in the **Operation** column.

Step 2 In the displayed **Pause Task** dialog box, select **Pause log capturing** and click **Yes**.

NOTE

- When a task in incremental state is paused, only the replay or capture and replay of incremental data is paused. Before database cutover, stop the task.
- After you select **Pause log capturing**, the DRS instance will no longer communicate with the source and destination databases. If the pause duration is too long, the task may fail to be resumed because the logs required by the source database expire. You are not advised pausing a task for more than 24 hours. For details, check the corresponding log configuration.
- After the task is paused, the status of the task becomes **Paused**.

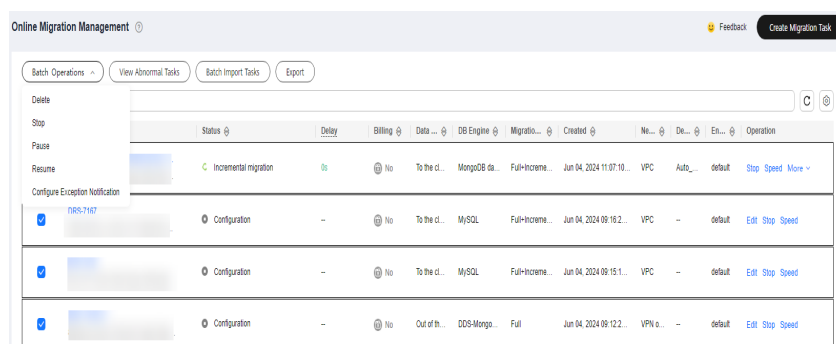
----End

Pausing Tasks

Step 1 On the **Online Migration Management** page, select the tasks to be paused.

Step 2 Click **Batch Operations** in the upper left corner and choose **Pause**.

Figure 4-43 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

4.6.10 Cloning a Migration Task

DRS allows you to quickly clone the configuration of an existing migration task. However, tasks in the following status cannot be cloned:

- Creating
- Creation failed
- Configuration
- Pending start
- Starting
- Deleted

You can clone the following migration tasks:

- To the cloud
 - MongoDB->DDS
- From the cloud
 - DDS->MongoDB

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.
- 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 migration task has been created.

Procedure

Step 1 On the **Online Migration 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

4.6.11 Stopping a Migration Task

After the source database and services are migrated to the destination database, you can stop the migration task. To prevent data from being overwritten after the source database and services are migrated to the destination database, operations on the source database should not be synchronized to the destination database. This section describes how to stop a migration task to achieve this goal.

You can stop a task in any of the following statuses:

- Creating
- Configuration
- Pending start
- Full migration
- Full migration failed
- Incremental migration
- Incremental migration 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 replication 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 resumed.
 - Migration tasks support only the pay-per-use billing. The billing stops after a task is stopped.
-

Prerequisites

- You have logged in to the DRS console.
- A migration task is in progress.

Stopping a Task

Step 1 On the **Online Migration Management** page, locate the task and click **Stop** in the **Operation** column.

Step 2 In the displayed dialog box, click **Yes**.

NOTE

- Generally, triggers and events will be synchronized when you stop the task.
- 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 and will not migrate triggers and events. You have to manually migrate triggers and events.
- If you need to migrate triggers and events, restore the DRS task first. After the task status becomes normal, you can stop the task.
- For a DRS task that is in the incremental state and with MySQL serving as the source database, after you select **Display breakpoint information when the task is stopped** when you stop the task, the GTID and binlog position information of the source database will be displayed on the migration progress page after the task is stopped.

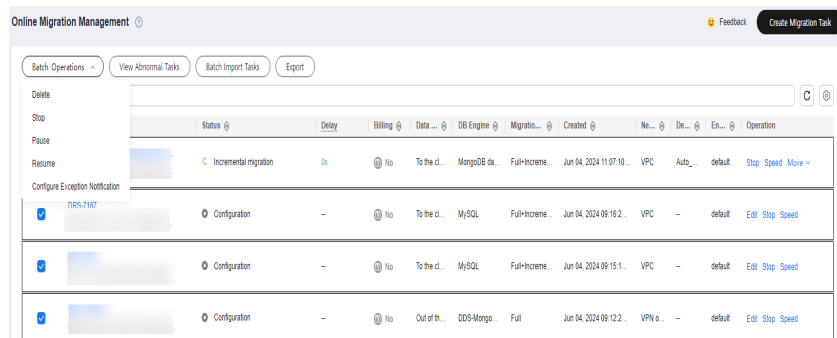
----End

Stopping Tasks

Step 1 In the task list on the **Online Migration Management** page, select the tasks you want to stop.

Step 2 Click **Batch Operations** in the upper left corner and choose **Stop**.

Figure 4-44 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

4.6.12 Deleting a Migration Task

This section describes how to delete a migration 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.
- A migration task that has been completed or fails to be configured exists.

Deleting a Task

Step 1 In the task list on the **Online Migration Management** page, locate the target task and click **Delete** in the **Operation** column.

Step 2 Click **Yes** to submit the deletion task.

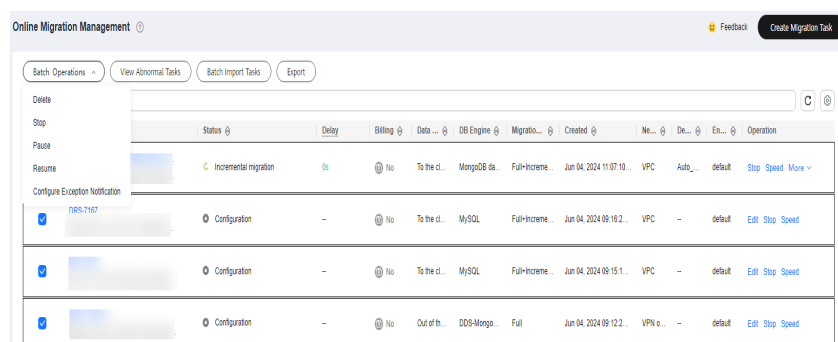
----End

Deleting Tasks

Step 1 On the **Online Migration Management** page, select the tasks to be deleted.

Step 2 Click **Batch Operations** in the upper left corner and choose **Delete**.

Figure 4-45 Batch Operations



Step 3 In the displayed dialog box, confirm the task information and click **Yes**.

----End

4.6.13 Importing Tasks in Batches

DRS allows you to import an Excel file to create migration tasks in batches. You need to fill in a template with basic task information, source and destination database information, and objects to be migrated, and then import the template.

Constraints

- Only MySQL-to-MySQL tasks 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 only set 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.

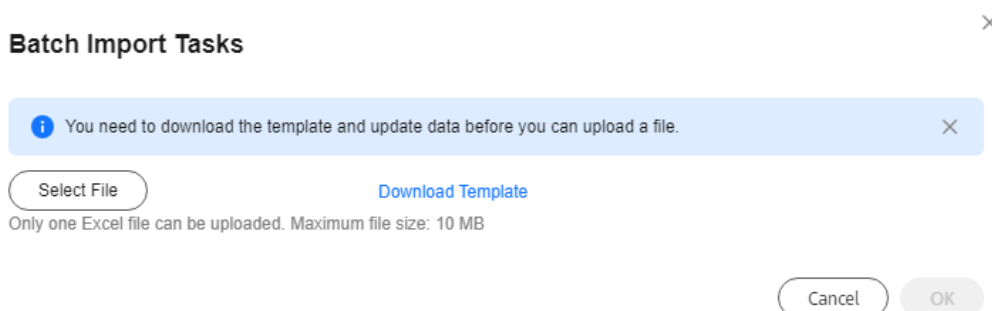
- You cannot schedule the start time of a task in the import template. After the pre-check for the task is successful, you can click **Edit** on the GUI to access the startup page and set a scheduled start time.
- The system cannot verify the correctness of all parameters in the import template. You 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 **Online Migration Management** page, click **Batch Import Tasks** in the upper left corner.

Step 2 On the **Batch Import Tasks** dialog box, click **Download Template**.

Figure 4-46 Batch Import Tasks



Step 3 Edit the downloaded Excel template and enter task information as required.

Step 4 Click **Select File** to upload the template.

Step 5 After the upload and verification are successful, click **Confirm** to create tasks in batches.

NOTE

- DRS does not save the uploaded Excel file and only parses the content in the file.
- The IP addresses, domain names, usernames, and passwords of the source and destination databases are encrypted and stored in the system until a task is deleted.

----End

4.6.14 Upgrading the Version of a DRS Task

In 24.10.0 and later versions, you can upgrade the version of a DRS task on the console immediately or as scheduled. You can set an upgrade time window. During the time window, the system checks whether a task meets the upgrade conditions every 10 minutes. If the task meets the upgrade conditions, the system delivers the request to the kernel for version upgrade.

Constraints

- The upgrade conditions displayed on the console are as follows:
 - The task is in the incremental state.
 - The kernel version must be 24.10.0 or later. To upgrade the DRS kernel, submit a service ticket by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console.
 - The task has a version that can be upgraded.
 - The version upgrade function is available only for synchronization, migration, and DR tasks.
- The upgrade conditions on which the management system depends for delivering an upgrade command are as follows:
 - The task is in the incremental state.
 - The task latency is no more than 30s.
 - The task has a version that can be upgraded, and the upgrade time is within a specified time range.
 - In the multi-task scenario, an upgrade can be performed only when all subtasks meet the preceding conditions.
- After receiving the upgrade request from the management system, the kernel downloads the new version and monitors and checks whether the upgrade is successful based on the following conditions:
 - The log download, parsing, and incremental migration processes are normal.
 - The incremental position is updated properly.
- The maximum monitoring duration is 10 minutes. (If the system detects that the incremental position is updated properly within 3 minutes and the preceding processes are normal within 3 minutes, the monitoring terminates within 3 minutes.)
- If the task upgrade fails, the system rolls back the version. The entire upgrade process is displayed in the **Migration Logs** page on the console.
- The task cannot be paused during version upgrade.

Prerequisites

- The task is in the incremental state. The task has been upgraded to the baseline version. The **Upgrade** button is available in the task list.

Procedure

- Step 1** On the **Online Migration Management** page, locate the target migration task and click **Upgrade** in the **Operation** column.

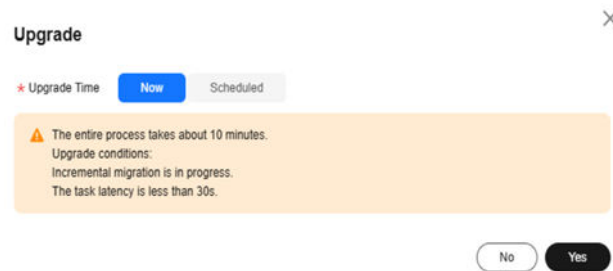
Figure 4-47 Task management

Task NameID	Status	Delay	Billing	Data Flow	DB Engine	Migration Type	Created	Retain...	Discrep...	Enterp...	Operation
DRS-4735-migrate-upgrade 348e12c1-2664-480e-b817-650aa@195	Incremental migration	Incremental delay: 3s 0 RPS (0 KPS)	no	To the cloud	DRS	Full-Incremental	Oct 26, 2024 15:28:52 GMT+08:00	VPC	default	Upgrade	Show

- Step 2** In the **Upgrade** dialog box, select **Now** or **Scheduled** for **Upgrade Time**.

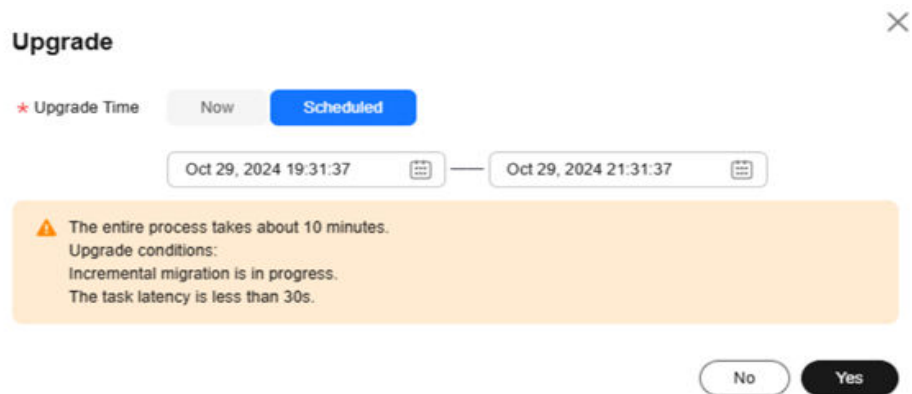
1. After you click **Now**, the task upgrade starts.

Figure 4-48 Upgrading a task



2. If you select **Scheduled** for **Upgrade Time**, specify a time range for task upgrade.

Figure 4-49 Upgrading a task as scheduled



----End

Viewing the Upgrade Status

- During the upgrade, the task status is **Upgrading replication instance** in the task list.

Figure 4-50 Upgrading

Task Name	Status	Detail	Billing	Data Flow	DB Engine	Migration Type	Created	Network	Default	Endpoint	Operation
DRS-4788-mongodb-upgrade	Incremental migration	Incremental migration	DRS	To the cloud	DDS	Full-Incremental	Oct 26, 2024 15:29:02 GMT+08:00	VPC	default		Upgrade, Show
3456789-2000-4567-8910-123456789101	Upgrading replication instance										

- Click the task name to go to the task details page. In the navigation pane, choose **Migration Logs** to view upgrade logs.
 - Upgrade success logs

Figure 4-51 Logs

< | DRS-4788-mongodb-upgrade

Time	Level	Description
Oct 26, 2024 10:58:36 GMT+08:00	Info	version upgrade from 24.10.0.0 to 24.10.0.2 completed, status successful
Oct 26, 2024 10:58:15 GMT+08:00	Info	increment transfer start
Oct 26, 2024 10:58:06 GMT+08:00	Info	version upgrade from 24.10.0.0 to 24.10.0.2 started

- Upgrade failure logs

Figure 4-52 Logs

Time	Level	Description
Oct 26, 2024 10:51:35 GMT+08:00	Info	version upgrade from 24.10.0.0 to 24.10.0.1 completed, status failed
Oct 26, 2024 10:51:35 GMT+08:00	Info	rollback to version 24.10.0.0 successful
Oct 26, 2024 10:51:23 GMT+08:00	Info	increment transfer start
Oct 26, 2024 10:51:07 GMT+08:00	Info	version upgrade from 24.10.0.0 to 24.10.0.1 failed, err/msg/task:CheckUpgradeStatusTask error, rollback to version 24.10.0.0 started
Oct 26, 2024 10:49:17 GMT+08:00	Info	increment transfer start
Oct 26, 2024 10:47:21 GMT+08:00	Info	version upgrade from 24.10.0.0 to 24.10.0.1 started

4.6.15 Task Statuses

Migration statuses indicate different migration phases.

Table 4-17 lists statuses and descriptions of online migration tasks.

Table 4-17 Task status and description

Status	Description
Creating	A replication instance is being created for DRS.
Task creation failed	Failed to create a replication instance for real-time migration.
Configuration	A replication instance is created, but the migration 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 migration task has been delivered to the replication instance, waiting for the replication instance to start the migration task.
Starting	A migration task is being started.
Start failed	Failed to start a real-time migration task.
Full migration	A full migration task is being performed.
Full migration failed	Failed to perform a full migration task.
Incremental migration	An incremental migration task is being performed.
Incremental migration failed	Failed to perform an incremental migration task.

Status	Description
Fault rectification	A replication instance is faulty and the system automatically restores the migration task.
Cloning	A migration task is being cloned.
Cloning failed	Failed to clone a migration task.
Paused	A real-time migration task is paused.
Stopping	The replication instance and resources used for executing the migration task are being released.
Completing	A replication instance and resources are being released.
Stopping task failed	Failed to release the replication instance and resources used by the migration task.
Completed	The task is completed and the replication 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 migration tasks are not displayed in the status list.

5 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.

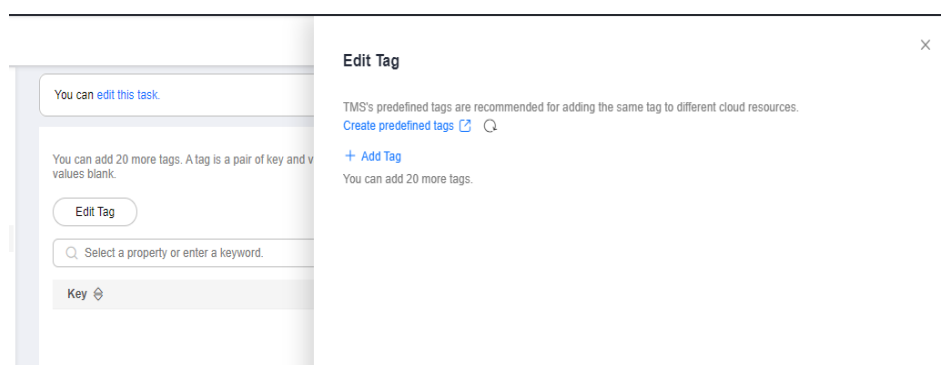
- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- Each DB instance can have up to 20 tags.

Adding a Tag

Step 1 On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 2 In the navigation pane on the left, choose **Tags**.

Step 3 On the **Tags** page, click **Edit Tag**. In the displayed dialog box, click **Add Tag**, enter a tag key and value, and click **OK**.



- When you enter a tag key and value, the system automatically displays all tags (including predefined tags and resource tags) associated with all DB instances except the current one.

- The tag key cannot be empty and must be unique. It cannot start or end with a space or start with **_sys_**. It can contain 1 to 128 characters, including letters, digits, spaces, and special characters **_:=+@**
- The tag value can be empty. It cannot start or end with a space and can contain 0 to 255 characters, including letters, digits, spaces, and special characters **_:=+@**

Step 4 View and manage the tag on the **Tags** page.

----End

Editing a Tag

Step 1 On the **Online Migration Management** page, click the target migration 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 **Online Migration Management** page, click the target migration 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

6 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 6-1 Quick Diagnosis

Source Database

System databases, accounts, and parameters will not be synchronized. You need to manually create accounts and configure parameters in parameter templates of the destination database.

IP Address or Domain Name

Port

Database Username

Database Password

SSL Connection

Test Connection

The network connection between the replication instance and the database is faulty. For a cross-VPC task, the network between different VPCs may be disconnected. For details about how to create a VPC peering connection, see the VPC documentation, [View details](#)

Quick Diagnosis

Step 3 View the diagnosis result on the displayed **Diagnosis Details** dialog box. The result includes the packet loss rate and port check result.

Figure 6-2 Diagnosis Details

Diagnosis Details ×

IP Address or Domain Name	Packet Loss Rate (%)	Port Check
<input type="text"/>	100	❌ Failed

OK

----End

7 Interconnecting with CTS

7.1 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 7-1 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

7.2 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

8 Interconnecting with Cloud Eye

8.1 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 8-1](#) lists the DRS performance metrics.

Table 8-1 DRS metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
cpu_util	CPU Usage	CPU usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
mem_util	Memory Usage	Memory usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
network_incoming_bytes_rate	Network Input Throughput	Incoming traffic in bytes per second	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
network_outgoing_bytes_rate	Network Output Throughput	Outgoing traffic in bytes per second	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
disk_read_bytes_rate	Disk Read Throughput	Number of bytes read from the disk per second (bytes/second).	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
disk_write_bytes_rate	Disk Write Throughput	Number of bytes written to the disk per second (bytes/second).	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
disk_util	Storage Space Usage	Storage space usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_bytes_rate	Source Database Read Throughput	Table data or WAL bytes read from the source database per second	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_rows_rate	Rows Read from Source Database per Second	Number of table data rows or WAL rows read from the source database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_latency	Source Database WAL Extract Lag	Latency of extracting WAL from the source database Unit: ms.	\geq ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_bytes_rate	Destination Database Write Throughput	Number of bytes written to the destination database per second.	≥ 0 byte/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_rows_rate	Rows Written into Destination Database per Second	Number of rows that are written to the destination database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_transactions_rate	DML TPS	Number of DML transactions written to the destination database per second.	≥ 0 transaction/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_ddls_number apply_ddls_rate NOTE apply_ddls_rate is replaced by apply_ddls_number after December 2022.	DDL TPS	Total number of DDL transactions written into the destination database.	≥ 0 transaction	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_latency	Replication Delay	Delay (in milliseconds) of data replay.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_average_execute_time	Average Transaction Execution Time	Average execution time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is millisecond.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_average_commit_time	Average Transaction Commit Time	Average commit time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is ms.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_current_state	Synchronization Status	This metric is the synchronization status of the current kernel data (10: abnormal; 1: idle; 2: DML; 3: DDL), instead of the task status.	10: abnormal 1: idle 2: DML is executed. 3: DDL is executed.	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_thread_workers	Synchronization Threads	Number of working threads for data synchronization	≥ 0	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_job_status	Task Status	Status of the current task. (0: normal; 1: abnormal; 2: paused)	0: normal 1: abnormal 2: paused	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Dimensions

Key	Value
instance_id	DRS instance ID


8.2 Configuring Alarm Rules

Scenarios

You can configure DRS alarm rules to customize the monitored objects and notification policies and learn the DRS running status in a timely manner.

This section describes how to set DRS alarm rules, including the alarm rule name, service, dimension, monitoring scope, template, and whether to send a notification.

Procedure

- Step 1** Log in to the management console.
- Step 2** Under **Management & Governance**, click **Cloud Eye**.
- Step 3** In the navigation pane on the left, choose **Cloud Eye > Data Replication Service**.
- Step 4** Select the DB instance which you want to create an alarm rule for and click **Create Alarm Rule** in the **Operation** column.
- Step 5** On the displayed page, set parameters as required.
- Specify **Name** and **Description**.
 - Select **Use template** for **Method**. The template contains the following common metrics: CPU usage, memory usage, and storage space usage.
 - Click  to enable alarm notification. The validity period is 24 hours by default. If the topics you required are not displayed in the drop-down list, click **Create an SMN topic**. Then, select **Generated alarm** and **Cleared alarm** for **Trigger Condition**.

NOTE

Cloud Eye sends notifications only within the validity period specified in the alarm rule.

- Step 6** Click **Create**. The alarm rule is created.

For details about how to create alarm rules, see [Creating an Alarm Rule](#) in the *Cloud Eye User Guide*.

----End

8.3 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

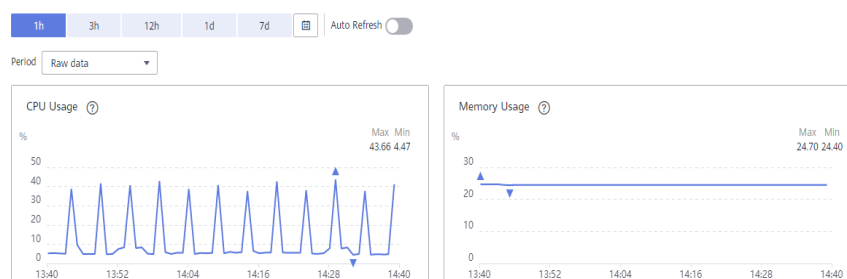
Viewing Metrics

- 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 8-1 Viewing monitoring metrics



----End

9 Interconnecting with LTS

9.1 Log Reporting

Scenarios

If you enable log reporting, all logs generated by DRS instances (including real-time migration, backup migration, real-time synchronization, real-time disaster recovery, and workload replay instances) are uploaded to Log Tank Service (LTS) for management.

Precautions

- After this function is enabled, all logs of the task are reported by default.
- This request does not take effect immediately. There is a delay of about 10 minutes.
- You will be billed for this function. For details, see [LTS Pricing Details](#).
- Ensure that there are available LTS log groups and log streams in the same region as your instance.

For more information about log groups and log streams, see [Log Management](#).

- After this function is disabled, you will not be billed anymore.

Enabling or Disabling Log Reporting

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and project.

Step 3 Choose **Database > Data Replication Service**. The **Data Replication Service** page is displayed.

Step 4 Take real-time migration as an example. On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column. The operations for real-time synchronization, real-time disaster recovery, and workload replay are similar to those for real-time migration.

Step 5 On the **Basic Information** page, click **Migration Logs** on the left.

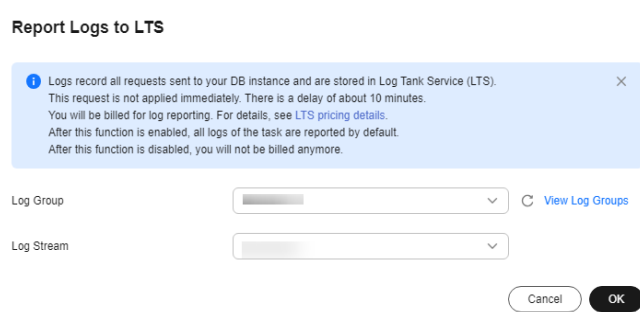
Step 6 Click  next to **Report Logs to LTS** in the upper part of the page.

Step 7 Select an LTS log group and log stream and click **OK**.

 **NOTE**

This request does not take effect immediately. There is a delay of about 10 minutes.


Figure 9-1 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 9-2 Disabling log reporting to LTS



----End

9.2 Viewing and Downloading Logs

Scenarios

If you have enabled log reporting to LTS for a DRS task in **Log Reporting**, you can analyze logs, search for logs, visualize logs, download logs, and view real-time logs on the LTS console.

Viewing Logs Reported to LTS


- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Under **Management & Governance**, click **Log Tank Service**.
- Step 4** In the **Log Groups** area, locate a target log group and click its name. For details about LTS, see [Log Tank Service \(LTS\) User Guide](#).

Figure 9-3 Viewing log details




Table 9-1 Log field description

Name	Type	Description
_resource_id	String	Resource ID. The value is fixed to projectId for DRS.
_resource_name	String	Resource name. The value is fixed to DRS .
_service_type	String	Service type. The value is fixed to Data Replication Service .

----End

Downloading Logs Reported to LTS



- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Under **Management & Governance**, click **Log Tank Service**.
- Step 4** In the **Log Groups** area, locate a target log group and click its name.
- Step 5** Click **Download** on the right to download logs. For details about LTS, see [Log Tank Service \(LTS\) User Guide](#).

Figure 9-4 Downloading logs



-----End

10 Appendix

10.1 General Usage Suggestions

Before using DRS, you need to read the following:

- The destination database version must be the same as or later than the source database version.
- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
- To maintain data consistency before and after the migration, do not write data to the source and destination databases in the full migration mode. In the full+incremental migration mode, you can continue the migration while data is still being written to the source database.
- When a task is being started or in the full migration 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 migrated without a primary key may be locked for 3s.
- The data being migrated may be locked by other transactions for a long period of time, resulting in read timeout.
- Due to the inherent characteristics of MySQL, in some 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 migration, 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?](#)

- It is best to start your migration task during off-peak hours by setting **Start Time** to **Start at a specified time**. A less active database is easier to migrate successfully and prevents impact on the database performance. If you have to migrate data during peak hours, you can select **Yes** for **Flow Control** to adjust the migration speed.

If many-to-one migration is involved, you need to understand the following:

- When creating multiple migration tasks in the many-to-one scenario, ensure that the read and write settings of the destination database are consistent in these tasks. For details, see [Constraints and Operation Suggestions on Many-to-One Scenario](#).

After data migration 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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

10.2 DRS Pre-Check Items

A DRS migration 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 migration, DRS automatically scans and checks parameters and objects in the source and destination databases before starting a migration task. The following table lists common check items.

Overview

View pre-check items based on the source database type of your migration task.

- The following lists the pre-check items for migration tasks with MySQL serving as the source.
 - [MySQL->MySQL](#)
 - [MySQL->DDM](#)
 - [MySQL->GaussDB\(for MySQL\)](#)
- The following lists the pre-check items for migration tasks with MySQL schema and logic table serving as the source.
 - [MySQL schema and logic table -> DDM](#)
- The following lists the pre-check items for migration tasks with MongoDB serving as the source.
 - [MongoDB->DDS](#)
 - [MongoDB->GeminiDB Mongo](#)
- The following lists the pre-check items for migration tasks with DDS serving as the source.
 - [DDS->MongoDB](#)

MySQL->MySQL

Table 10-1 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full migration requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT Full+incremental migration requires the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT The user must have the SELECT permission for mysql.user if the source database is a non-Alibaba Cloud database. If the source database is an Alibaba Cloud database, the user must have the SELECT permission for mysql.user and mysql.user_view. 	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> The destination database user must have the following permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database. 	For details, see Which MySQL Permissions Are Required for DRS?
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.
	Destination database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Migration version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version.
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database.
	Performance parameters	Enable skip-name-resolve for the source database to reduce the possibility of connection timeout.	-
		The log_slave_updates parameter of the source database must be enabled. Otherwise, the migration fails.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured.
		The binlog_row_image parameter value of the source database must be FULL . Otherwise, the migration fails.	For details, see Checking Whether the binlog_row_image Value is FULL.
	Maximum size of a packet that can be transmitted	If there is a lot of data to be migrated or there are too many fields to be migrated, and the max_allowed_packet value of the source database is too small, the migration task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small.
If there is a lot of data to be migrated or there are too many fields to be migrated, and the max_allowed_packet value of the destination database is too small, data cannot be written to the destination database and the full migration fails.		For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small.	

Category	Check Item	Check Item Details	Solution to Failure
	sql_mode value	If the MyISAM tables are included in the migration objects, the sql_mode value in the destination database cannot be no_engine_substitution . Otherwise, the migration fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental migration	Binlog status	The binlog of the source database must be enabled and the row-based format must be used during incremental migration.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	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 migration may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental migration, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .

Category	Check Item	Check Item Details	Solution to Failure
	Session variables	During an incremental migration, if the session variable character_set_client is set to binary , some data may include garbled characters.	-
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Transaction isolation level	The transaction isolation level of the destination database must be the same as that of the source database.	For details, see Checking Whether the Transaction Isolation Levels are Consistent.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	server_uuid parameter	The server_uuid values of the destination and source databases must be different.	For details, see Checking Whether the SERVER_UUID Values of the Source and Destination Databases Are the Same.
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Migration objects	Selected objects	<ul style="list-style-type: none"> • Databases, tables, users, views, indexes, constraints, functions, stored procedures, triggers, and events can be migrated. • Only MyISAM and InnoDB tables can be migrated. • The system database and event statuses cannot be migrated. 	-
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	For details, see Checking Whether Referenced Tables Are Selected for Migration.
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.

Category	Check Item	Check Item Details	Solution to Failure
	Duplicate names	The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same .
	Table naming rules	<ul style="list-style-type: none"> The source database name, table name, and view name cannot contain non-ASCII characters or special characters '<>`\'\"' The source database name cannot start with ib_logfile and cannot be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-
	Encrypted table check	If there are encrypted tables in the source database, check whether these tables are supported on the destination database. If the destination database does not support these tables, the task may fail.	For details, see Source Encrypted Table Check .

MySQL->DDM

Table 10-2 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full migration requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT Full+incremental migration requires the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT 	For details, see Which MySQL Permissions Are Required for DRS?

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> 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 migrated. 	For details, see Creating an Account .
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases .
	Destination database version	DDM instance version.	For details, see Supported Databases .
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	Enable skip-name-resolve for the source database to reduce the possibility of connection timeout.	-
		The log_slave_updates parameter of the source database must be enabled. Otherwise, the migration fails.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured .
		The binlog_row_image parameter value of the source database must be FULL . Otherwise, the migration fails.	For details, see Checking Whether the binlog_row_image Value is FULL .

Category	Check Item	Check Item Details	Solution to Failure
	Maximum size of a packet that can be transmitted	If there is a lot of data to be migrated or there are too many fields to be migrated, and the <code>max_allowed_packet</code> value of the source database is too small, the migration task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small .
		If there is a lot of data to be migrated or there are too many fields to be migrated, and the <code>max_allowed_packet</code> value of the destination database is too small, data cannot be written to the destination database and the full migration fails.	For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small .
	<code>sql_mode</code> value	If the MyISAM tables are included in the migration objects, the <code>sql_mode</code> value in the destination database cannot be <code>no_engine_substitution</code> . Otherwise, the migration fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental migration	Binlog status	The binlog of the source database must be enabled and the row-based format must be used during incremental migration.	For details, see Checking Whether the Source Database Binlog Is Enabled .

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 migration may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental migration, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .
	Session variables	<p>During an incremental migration, if the session variable character_set_client is set to binary, some data may include garbled characters.</p>	-
Destination database	Storage space	<p>The destination DB instance must have sufficient storage space.</p>	For details, see Checking Whether the Destination Database Has Sufficient Storage Space .

Category	Check Item	Check Item Details	Solution to Failure
	Status	The destination DB instance is running properly.	-
	auto_increment parameter	The auto_increment value of a table in the destination cannot be less than that of a table in the source.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.

Category	Check Item	Check Item Details	Solution to Failure
	Transaction isolation level	The transaction isolation level of the destination database must be the same as that of the source database.	For details, see Checking Whether the Transaction Isolation Levels are Consistent .
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent .
	server_uuid parameter	The server_uuid values of the destination and source databases must be different.	For details, see Checking Whether the SERVER_UUID Values of the Source and Destination Databases Are the Same .
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same .
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-

Category	Check Item	Check Item Details	Solution to Failure
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
	Table structures	<ul style="list-style-type: none"> To migrate table structures and other objects, you need to create table structures and indexes in the destination database based on table structures of the source schema. Any source objects that have no corresponding objects created in the destination cannot be migrated. The table structure created in the destination database must be the same as that in the source database. 	-
Migration objects	Selected objects	<ul style="list-style-type: none"> Only the source database data can be migrated. The table structure and other database objects of the source database cannot be migrated. Only MyISAM and InnoDB tables can be migrated. The system database and event statuses cannot be migrated. If the source database is MySQL 8.0, tables with function indexes cannot be migrated. 	-
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.

Category	Check Item	Check Item Details	Solution to Failure
	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	For details, see Checking Whether Referenced Tables Are Selected for Migration .
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database .
	Table naming rules	<ul style="list-style-type: none"> The names of the source databases and tables cannot contain non-ASCII characters, or special characters <'>`\" The source database name cannot start with ib_logfile and cannot be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-
	Encrypted table check	If there are encrypted tables in the source database, check whether these tables are supported on the destination database. If the destination database does not support these tables, the task may fail.	For details, see Source Encrypted Table Check .

MySQL->GaussDB(for MySQL)

Table 10-3 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full migration requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT Full+incremental migration requires the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT The user must have the SELECT permission for mysql.user if the source database is a non-Alibaba Cloud database. If the source database is an Alibaba Cloud database, the user must have the SELECT permission for mysql.user and mysql.user_view. 	For details, see Which MySQL Permissions Are Required for DRS?
	Destination database permissions	<ul style="list-style-type: none"> The destination database user must have the following permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database. 	For details, see Which MySQL Permissions Are Required for DRS?
Versions	Source database versions	Versions 5.5, 5.6, 5.7, and 8.0 are supported.	For details, see Supported Databases.
	Destination database version	GaussDB(for MySQL) 8.0 is supported.	For details, see Supported Databases.

Category	Check Item	Check Item Details	Solution to Failure
	Migration version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version.
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database.
	Performance parameters	Enable skip-name-resolve for the source database to reduce the possibility of connection timeout.	-
		The log_slave_updates parameter of the source database must be enabled. Otherwise, the migration fails.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured.
		The binlog_row_image parameter value of the source database must be FULL . Otherwise, the migration fails.	For details, see Checking Whether the binlog_row_image Value is FULL.
	Maximum size of a packet that can be transmitted	If there is a lot of data to be migrated or there are too many fields to be migrated, and the max_allowed_packet value of the source database is too small, the migration task may fail.	For details, see Whether the max_allowed_packet Value of the Source Database Is Too Small.
If there is a lot of data to be migrated or there are too many fields to be migrated, and the max_allowed_packet value of the destination database is too small, data cannot be written to the destination database and the full migration fails.		For details, see Checking Whether the max_allowed_packet Value of the Destination Database Is too Small.	

Category	Check Item	Check Item Details	Solution to Failure
	sql_mode value	If the MyISAM tables are included in the migration objects, the sql_mode value in the destination database cannot be no_engine_substitution . Otherwise, the migration fails.	For details, see Checking Whether the Source Database Contains Invalid sql_mode Values .
Incremental migration	Binlog status	The binlog of the source database must be enabled and the row-based format must be used during incremental migration.	For details, see Checking Whether the Source Database Binlog Is Enabled .
	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 migration may fail.</p> <ul style="list-style-type: none"> • If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental migration, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> • If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. • If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .

Category	Check Item	Check Item Details	Solution to Failure
	Session variables	During an incremental migration, if the session variable character_set_client is set to binary , some data may include garbled characters.	-
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly.	-
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Time zone	The time_zone value of the destination database must be the same as that of the source database.	For details, see Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Case sensitive	The lower_case_table_names value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.
	Transaction isolation level	The transaction isolation level of the destination database must be the same as that of the source database.	For details, see Checking Whether the Transaction Isolation Levels are Consistent.
	Maximum length of the calculation result of the group_concat function	The group_concat_max_len value of the destination database must be the same as that of the source database.	For details, see Checking Whether the Values of group_concat_max_len Are Consistent.
	server_uuid parameter	The server_uuid values of the destination and source databases must be different.	For details, see Checking Whether the SERVER_UUID Values of the Source and Destination Databases Are the Same.
	InnoDB check mode	The innodb_strict_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.

Category	Check Item	Check Item Details	Solution to Failure
	Data block encryption parameter	The block_encryption_mode value of the destination database must be the same as that of the source database.	-
	SQL mode	The sql_mode value of the destination database must be the same as that of the source database.	For details, see Checking Whether the sql_mode Values of the Source and Destination Databases Are the Same.
Migration objects	Selected objects	<ul style="list-style-type: none"> • Databases, tables, users, views, indexes, constraints, functions, stored procedures, triggers, and events can be migrated. • Only MyISAM and InnoDB tables can be migrated. • The system database and event statuses cannot be migrated. 	-
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys.
	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	For details, see Checking Whether Referenced Tables Are Selected for Migration.
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database.

Category	Check Item	Check Item Details	Solution to Failure
	Duplicate names	The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).	For details, see Checking Whether the Names of the Source and Destination Databases Are the Same .
	Table naming rules	<ul style="list-style-type: none"> The source database name, table name, and view name cannot contain non-ASCII characters or special characters '<>`\'\"' The source database name cannot start with ib_logfile and cannot be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-
	Encrypted table check	If there are encrypted tables in the source database, check whether these tables are supported on the destination database. If the destination database does not support these tables, the task may fail.	For details, see Source Encrypted Table Check .

MySQL schema and logic table -> DDM

Table 10-4 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> Full migration requires the following minimum permissions: SELECT, SHOW VIEW, and EVENT Full+incremental migration requires the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT 	For details, see Which MySQL Permissions Are Required for DRS?

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	<ul style="list-style-type: none"> 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 migrated. 	For details, see Creating an Account .
Versions	Source database version	DDM instance version.	For details, see Supported Databases .
	Destination database version	DDM instance version.	For details, see Supported Databases .
Parameters	GTID status	Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.	For details, see Checking Whether GTID Is Enabled for the Source Database .
	Performance parameters	Enable skip-name-resolve for the source database to reduce the possibility of connection timeout.	-
		The log_slave_updates parameter of the source database must be enabled. Otherwise, the migration fails.	For details, see Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured .
		The binlog_row_image parameter value of the source database must be FULL . Otherwise, the migration fails.	For details, see Checking Whether the binlog_row_image Value is FULL .
Incremental migration	Binlog status	The binlog of the source database must be enabled and the row-based format must be used during incremental migration.	For details, see Checking Whether the Source Database Binlog Is Enabled .

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 migration may fail.</p> <ul style="list-style-type: none"> If the source database is a self-managed MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. If the source database is RDS for MySQL, set the binlog retention period by following the instructions provided in Setting a Local Retention Period for RDS for MySQL Binlogs. 	For details, see Checking Whether the Binlog Retention Period Is Set on the Source Database .
	server_id value	<p>During an incremental migration, the server_id value of the MySQL source database must be set.</p> <ul style="list-style-type: none"> If the source database version is MySQL 5.6 or earlier, the value of server_id ranges from 2 to 4294967296. If the source database version is MySQL 5.7 or later, the value of server_id ranges from 1 to 4294967296. 	For details, see Checking Whether the Source Database server_id Meets the Incremental Migration Requirements .
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space .
	Status	The destination DB instance is running properly.	-
	auto_increment parameter	The auto_increment value of a table in the destination cannot be less than that of a table in the source.	-

Category	Check Item	Check Item Details	Solution to Failure
Consistency	Character set	The character set of the destination database must be the same as that of the source database.	For details, see Checking Whether the Source and Destination Database Character Sets Are Consistent.
	Collation	The collation_server value of the destination database must be the same as that of the source database.	For details, see Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.
	Table structures	<ul style="list-style-type: none"> To migrate table structures and other objects, you need to create table structures and indexes in the destination database based on table structures of the source schema. Any source objects that have no corresponding objects created in the destination cannot be migrated. The table structure created in the destination database must be the same as that in the source database. 	-
Migration objects	Selected objects	<ul style="list-style-type: none"> Only the source database data can be migrated. The table structure and other database objects of the source database cannot be migrated. Only MyISAM and InnoDB tables can be migrated. The system database and event statuses cannot be migrated. If the source database is MySQL 8.0, tables with function indexes cannot be migrated. 	-

Category	Check Item	Check Item Details	Solution to Failure
	Tables without primary keys	Create primary keys for tables as the performance of a table without a primary key is lower than that of a table with a primary key.	For details, see Checking Whether the Source Database Tables Contain Primary Keys .
	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	For details, see Checking Whether Referenced Tables Are Selected for Migration .
	Foreign key reference operations	Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.	For details, see Whether There Are Foreign Keys Containing Unsupported Reference Operations in the Source Database .
	Destination database objects	Ensure that the destination database is empty before starting the migration. Otherwise, data in the destination may be overwritten during incremental migration.	-
	Table naming rules	<ul style="list-style-type: none"> • The names of the source databases and tables cannot contain non-ASCII characters, or special characters <'>`\" • The source database name cannot start with ib_logfile and cannot be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1. 	-

MongoDB->DDS

Table 10-5 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> ● Full migration requires the following minimum permissions: <ul style="list-style-type: none"> – Replica set: The source database user must have the readAnyDatabase permission for the admin database. – Single node: The source database user must have the readAnyDatabase permission for the admin database. – Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. – To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database. ● Full+incremental migration requires the following minimum permissions: <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 mongos node user must have the readAnyDatabase permission for the admin database and the 	For details, see MongoDB Migration and Synchronization .

Category	Check Item	Check Item Details	Solution to Failure
		<p>read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database and the read permission for the local database.</p> <ul style="list-style-type: none"> - To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database. 	
	Destination database permissions	<p>The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database.</p> <p>If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.</p>	For details, see DDS Migration and Synchronization .
Versions	Source database version	Versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4 are supported.	For details, see Supported Databases .
	Destination database version	Versions 3.4, 4.0, 4.2, and 4.4 are supported.	For details, see Supported Databases .
	Migration version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .
Incremental migration	Oplog status	During an incremental migration, Oplog of the source database must be enabled.	-
	Oplog retention period	If the storage space is sufficient, store the source database Oplog for as long as possible. The recommended retention period is three days.	-

Category	Check Item	Check Item Details	Solution to Failure
	Balancer status	If the incremental source data cluster is migrated, the source database balancer must be disabled.	For details, see How Do I Disable the Balancer?
	Orphaned documents	If the full+incremental data needs to be migrated from the source cluster database, the balancer of the source database must be disabled and orphaned documents must be deleted.	For details, see How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
		DRS supports full migration between cluster instances. If the source cluster instance is not sharded, ensure that the size of the primary shard on the destination database is greater than that of the source database.	
	Status	The destination DB instance is running properly.	-
	Maximum number of chunks	Check whether the maximum number of chunks in the destination database is sufficient to support sharding and splitting of the source database. If the maximum number of chunks is reached, chunks are not split and the write performance is negatively affected.	For details, see Checking Whether the Maximum Number of Chunks in the Destination Database Is Sufficient.
Consistency	Collections	Check whether collections in both the source and destination databases are not capped. If not, the migration fails.	For details, see Checking Whether Collections in Both the Source and Destination Databases Are Not Capped.
Migration objects	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	-

Category	Check Item	Check Item Details	Solution to Failure
	Referenced roles	If the roles referenced by accounts to be migrated are not migrated to the destination database, the migration may fail.	For details, see Checking Whether the Source Database Referenced Roles Pass the Check.
	Referenced accounts	If the roles referenced by accounts to be migrated are not migrated to the destination database, the migration may fail.	For details, see Checking Whether the Source Database Referenced Accounts Pass the Check.
	Duplicate names	Destination database cannot contain non-empty collections with the same name as those in the source database.	For details, see Checking Whether the Destination Database Contains a Non-Empty Collection with the Same Name As the Source Database.
		Destination database cannot contain views with the same name as those in the source database.	For details, see Checking Whether the Same View Names Exist in Both the Source and Destination Databases.
	Naming rules	The source database name cannot contain \, "\$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).	-
	Composite hash indexes	<ul style="list-style-type: none"> 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. 	For details, see Whether There Are Composite Hash Indexes in the Source Collection.

Category	Check Item	Check Item Details	Solution to Failure
	Composite hash shard keys		For details, see Whether There Are Composite Hash Shard Keys in the Source Collection.
	Indexes	Check whether the <code>_id</code> fields in the collections of the source database have indexes. If not, the migration fails.	For details, see Checking Whether the <code>_id</code> Fields in the Collection of the Source Database Have Indexes.
	TTL indexes	Due to inconsistency of time zones and clocks between source and destination databases, migrating TTL indexes will cause data inconsistency.	For details, see Checking Whether the Source Database Collections Contain TTL Indexes.
	The number of indexes	If the number of indexes in the source database exceeds 10, the migration duration is affected.	For details, see Checking Whether the Source Database Collections Contain More Than 10 Indexes.

MongoDB->GeminiDB Mongo

Table 10-6 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> ● Full migration requires the following minimum permissions: <ul style="list-style-type: none"> – Replica set: The source database user must have the readAnyDatabase permission for the admin database. – Single node: The source database user must have the readAnyDatabase permission for the admin database. – Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. – To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database. ● Full+incremental migration requires the following minimum permissions: <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 mongos node user must have the readAnyDatabase permission for the admin database and the 	For details, see MongoDB Migration and Synchronization .

Category	Check Item	Check Item Details	Solution to Failure
		<p>read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database and the read permission for the local database.</p> <ul style="list-style-type: none"> - To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the system.users and system.roles system tables of the admin database. 	
	Destination database permissions	<p>The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database.</p> <p>If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.</p>	For details, see MongoDB Migration and Synchronization .
Versions	Source database version	Versions 3.2, 3.4, 4.0, and 4.2 are supported.	For details, see Supported Databases .
	Destination database version	Versions 3.4 and 4.0 are supported.	For details, see Supported Databases .
	Migration version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .
Incremental migration	Oplog status	During an incremental migration, Oplog of the source database must be enabled.	-
	Oplog retention period	If the storage space is sufficient, store the source database Oplog for as long as possible. The recommended retention period is three days.	-

Category	Check Item	Check Item Details	Solution to Failure
	Balancer status	If the incremental source data cluster is migrated, the source database balancer must be disabled.	For details, see How Do I Disable the Balancer?
	Orphaned documents	If the full+incremental data needs to be migrated from the source cluster database, the balancer of the source database must be disabled and orphaned documents must be deleted.	For details, see How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
		DRS supports full migration between cluster instances. If the source cluster instance is not sharded, ensure that the size of the primary shard on the destination database is greater than that of the source database.	
	Status	The destination DB instance is running properly.	-
	Maximum number of chunks	Check whether the maximum number of chunks in the destination database is sufficient to support sharding and splitting of the source database. If the maximum number of chunks is reached, chunks are not split and the write performance is negatively affected.	For details, see Checking Whether the Maximum Number of Chunks in the Destination Database Is Sufficient.
Consistency	Collections	Check whether collections in both the source and destination databases are not capped. If not, the migration fails.	For details, see Checking Whether Collections in Both the Source and Destination Databases Are Not Capped.
Migration objects	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	-

Category	Check Item	Check Item Details	Solution to Failure
	Referenced roles	If the roles referenced by accounts to be migrated are not migrated to the destination database, the migration may fail.	For details, see Checking Whether the Source Database Referenced Roles Pass the Check.
	Referenced accounts	If the roles referenced by accounts to be migrated are not migrated to the destination database, the migration may fail.	For details, see Checking Whether the Source Database Referenced Accounts Pass the Check.
	Duplicate names	Destination database cannot contain non-empty collections with the same name as those in the source database.	For details, see Checking Whether the Destination Database Contains a Non-Empty Collection with the Same Name As the Source Database.
		Destination database cannot contain views with the same name as those in the source database.	For details, see Checking Whether the Same View Names Exist in Both the Source and Destination Databases.
	Naming rules	The source database name cannot contain \. "\$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).	-
	Indexes	Check whether the _id fields in the collections of the source database have indexes. If not, the migration fails.	For details, see Checking Whether the _Id Fields in the Collection of the Source Database Have Indexes.

Category	Check Item	Check Item Details	Solution to Failure
	TTL indexes	Due to inconsistency of time zones and clocks between source and destination databases, migrating TTL indexes will cause data inconsistency.	For details, see Checking Whether the Source Database Collections Contain TTL Indexes.
	The number of indexes	If the number of indexes in the source database exceeds 10, the migration duration is affected.	For details, see Checking Whether the Source Database Collections Contain More Than 10 Indexes.

DDS->MongoDB

Table 10-7 Pre-check items

Category	Check Item	Check Item Details	Solution to Failure
Permissions	Source database permissions	<ul style="list-style-type: none"> ● Full migration requires the following minimum permissions: <ul style="list-style-type: none"> – Replica set: The source database user must have the readAnyDatabase permission for the admin database. – Single node: The source database user must have the readAnyDatabase permission for the admin database. – Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. ● Full+incremental migration requires the following minimum permissions: <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 mongos node user must have the readAnyDatabase permission for the admin database and the read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database and the read permission for the local database. 	For details, see DDS Migration and Synchronization .

Category	Check Item	Check Item Details	Solution to Failure
	Destination database permissions	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 clusterManager permission for the admin database.	For details, see MongoDB Migration and Synchronization .
Versions	Source database version	Versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4 are supported.	For details, see Supported Databases .
	Destination database version	Versions 3.2, 3.4, 3.6, 4.0, 4.2, and 4.4 are supported.	For details, see Supported Databases .
	Migration version	The destination database version must be the same as or later than the source database version.	For details, see Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version .
Incremental migration	Olog status	During an incremental migration, Olog of the source database must be enabled.	-
	Olog retention period	If the storage space is sufficient, store the source database Olog for as long as possible. The recommended retention period is three days.	-
	Balancer status	If the incremental source data cluster is migrated, the source database balancer must be disabled.	For details, see How Do I Disable the Balancer?
	Orphaned documents	If the full+incremental data needs to be migrated from the source cluster database, the balancer of the source database must be disabled and orphaned documents must be deleted.	For details, see How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?

Category	Check Item	Check Item Details	Solution to Failure
Destination database	Storage space	The destination DB instance must have sufficient storage space.	For details, see Checking Whether the Destination Database Has Sufficient Storage Space.
	Status	The destination DB instance is running properly.	-
	Maximum number of chunks	Check whether the maximum number of chunks in the destination database is sufficient to support sharding and splitting of the source database. If the maximum number of chunks is reached, chunks are not split and the write performance is negatively affected.	For details, see Checking Whether the Maximum Number of Chunks in the Destination Database Is Sufficient.
Consistency	Collections	Check whether collections in both the source and destination databases are not capped. If not, the migration fails.	For details, see Checking Whether Collections in Both the Source and Destination Databases Are Not Capped.
Migration objects	Associated objects	Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.	-
	Duplicate names	Destination database cannot contain non-empty collections with the same name as those in the source database.	For details, see Checking Whether the Destination Database Contains a Non-Empty Collection with the Same Name As the Source Database.

Category	Check Item	Check Item Details	Solution to Failure
		Destination database cannot contain views with the same name as those in the source database.	For details, see Checking Whether the Same View Names Exist in Both the Source and Destination Databases.
	Naming rules	The source database name cannot contain \. "\$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).	-
	The number of indexes	If the number of indexes in the source database exceeds 10, the migration duration is affected.	For details, see Checking Whether the Source Database Collections Contain More Than 10 Indexes.