Data Replication Service

Real-Time Migration

Issue	01
Date	2022-09-30

HUAWEI TECHNOLOGIES CO., LTD.

Huawei Technologies Co., Ltd.

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

Website: https://www.huawei.com

Email: support@huawei.com

Contents

1 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
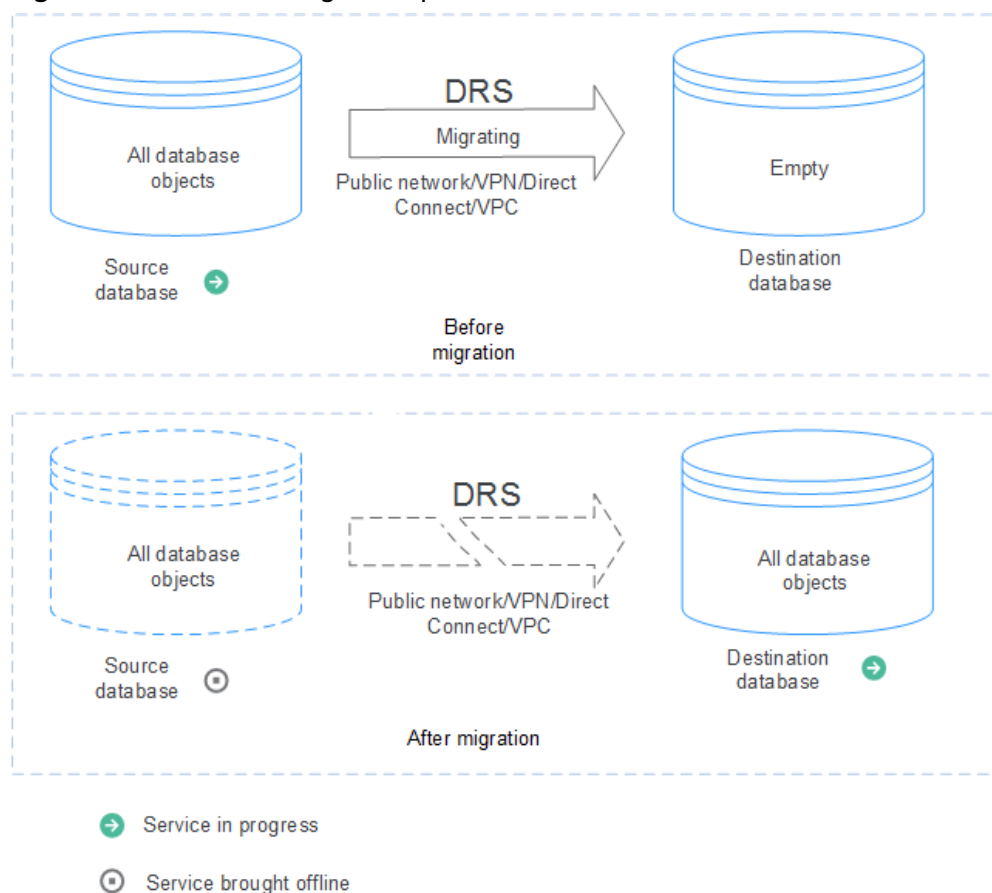
Supported Database Types

Table 1-1 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
• 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) primary/ standby	Full Full+Incremental	**From MySQL to GaussDB(for MySQL) Primary/ Standby**
• RDS for MySQL	• 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)
• On-premises MongoDB databases • MongoDB databases on an ECS • MongoDB database on other clouds • DDS	DDS	Full Full+Incremental	**From MongoDB to DDS**
• On-premises MongoDB databases • MongoDB databases on an ECS • MongoDB database on other clouds	GaussDB(for Mongo)	Full Full+Incremental: • Replica set -> Replica set • Replica set -> Cluster • Cluster -> Cluster	**From MongoDB to GaussDB(for Mongo)**

Source DB	Destination DB	Migration Type	Documentation
● DDS	● On-premises MongoDB databases ● MongoDB databases on an ECS ● MongoDB database on other clouds	Full Full+Incremental	**From DDS to MongoDB**
● On-premises MyCAT middleware ● MyCAT middleware on an ECS ● DDM	DDM	Full Full+Incremental	**From MySQL Schema and Logic Table to DDM**

2 To the Cloud

2.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 2-1 Supported databases

Source DB	Destination DB
● On-premises MySQL databases ● MySQL databases on an ECS ● MySQL databases on other clouds ● RDS for MySQL	● RDS for MySQL

Prerequisites

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

Suggestions

> ⚠️ **CAUTION**
>
> - When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
> - 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.

- 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.

- 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 you have to migrate data during peak hours, you can select **Yes** for **Flow Control** to adjust the migration speed.

  - 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.

  - 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%.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-level comparison

  To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

Precautions

Before creating a migration task, read the following notes.

Table 2-2 Precautions

Type	Restrictions
Database permissions	● Full migration (minimum permissions): – The source database user must have the following permissions: SELECT, SHOW VIEW, and EVENT – 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. ● Full+incremental migration (minimum permissions): – The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT – The destination database user must have the following permissions: SELECT, CREATE, 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. ● Account migration (minimum permissions): – 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. – The destination database user must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database.

Type	Restrictions
Source database	- The source database names cannot contain non-ASCII characters or special characters '<`>/\" - The names of the source tables and views cannot contain non-ASCII characters or special characters '<>/\" - The source database name cannot be **ib_logfile**. - The binlog of the MySQL source database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. - If the **expire_logs_days** value of the source database is **0**, the migration may fail. - During an incremental migration, the **server_id** value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the **server_id** value ranges from **2** to **4294967296**. If the source database is MySQL 5.7 or later, the **server_id** value ranges from **1** to **4294967296**. - Enable **skip-name-resolve** for the MySQL source database to reduce the possibility of connection timeout. - Enable the Global Transaction Identifier (GTID) of the source database. - The source database does not support the **mysql binlog dump** command. - The character sets of the source and destination databases must be the same. Otherwise, the migration fails. - The **log_slave_updates** parameter of the source database must be enabled. Otherwise, the migration fails. - The **binlog_row_image** parameter value of the source database must be **FULL**. Otherwise, the migration fails. - If the value of **lower_case_table_names** of the MySQL 8.0 source database is **0**, the database cannot be migrated.

Type	Restrictions
Destination database	• Data cannot be migrated from a newer version database to an older version database. • You are advised to use the row-based binlog in the destination MySQL database. Otherwise, an error may occur during an incremental migration. • The destination DB instance is running properly. • The destination DB instance must have sufficient storage space. • The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database). • The destination database isolation level must be set to at least read committed. • During migration, a large amount of data is written to the destination database. If the value of the **max_allowed_packet** parameter of the destination database is too small, data cannot be written. You are advised to set the **max_allowed_packet** parameter to a value greater than 100 MB. • Enable GTID of the destination database. • If the **server_uuid** values of the source and destination databases are the same, the incremental migration fails. • The **collation_server** values of the source and destination databases must be the same. Otherwise, the migration fails. • The table containing the foreign key must be migrated with the referenced tables. Otherwise, the migration fails. • The **time_zone** values of the source and destination databases must be the same. Otherwise, the migration fails. • The **sql_mode** values of the source and destination databases must be the same. Otherwise, the migration fails. • If the MyISAM tables are included in the migration objects, the **sql_mode** parameter in the destination database cannot contain the **no_engine_substitution** parameter. Otherwise, the migration fails. • The **innodb_strict_mode** values of the source and destination databases must be the same. Otherwise, the migration fails. • The **lower_case_table_names** values of the source and destination databases must be the same. Otherwise, the migration fails. • The **log_bin_trust_function_creators** parameter value of the destination database must be set to **on**. Otherwise, the migration fails.

Type	Restrictions
Migration objects	● Supported objects: databases, tables, views, indexes, constraints, functions, stored procedures, triggers, and events. ● The system database and event statuses cannot be migrated. ● Tables with storage engine different to MyISAM and InnoDB tables cannot be migrated.

Type	Restrictions
Precautions	● Objects that have dependencies must be migrated at the same time to avoid migration failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys ● Cascade operations cannot be performed on tables with foreign keys. ● 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. ● If you create 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. ● The destination database cannot be restored to a point in time when a full migration was being performed. ● If the source and destination sides are RDS MySQL instances, transparent data encryption (TDE) is not supported, and tables with the encryption function cannot be created. ● 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 and 5.7.28), you need to submit an O&M application for testing the SSL connection. ● The source database cannot be RDS read replicas of Alibaba Cloud. ● 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 destination database of a migration task can be set to **Read-only** or **Read/Write**. – **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 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.

Type	Restrictions
	• Set the **expire_log_day** parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • During task startup or full migration, DDL operations, such as deleting databases, indexes, and views, may cause the migration task to fail. • During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL and DML operations) during migration. • 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** on the source side. • 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 operating system or the database breaks down. • DDL statements are supported in full migration. • If the source and destination databases are of the same major version and the entire instance is migrated, DCL statements can be migrated in incremental mode, but users cannot be changed by updating the **mysql.user** table. For details about DCL statements, see the **MySQL official document**. • 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.

Procedure

This section uses the migration from MySQL to RDS 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, configure task details, description, and the replication instance, and click **Next**.

Figure 2-1 Migration task information

Table 2-3 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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 2-2 Replication instance information

Table 2-4 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**. Available options: **VPC**, **VPN or Direct Connect**, and **Public network**. By default, the value is **Public network**. ● VPC is suitable for migrations of cloud databases. ● Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP. ● VPN is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. ● Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. For details about networks, see **Preparations**.
Destination DB Instance	The RDS DB instance you created.
Replication Instance Subnet	The subnet where the replication instance resides. You can also click **View Subnet** to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Destination Database Access	● **Read-only** During 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 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. The task cannot be modified after being created.
Migration Type	● **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.
Enterprise Project	● If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. ● You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.

Parameter	Description
Tags	• This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. • After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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**.

📖 NOTE

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

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

Figure 2-3 Self-build on ECS - source database information

Table 2-5 Self-build on ECS - source database information

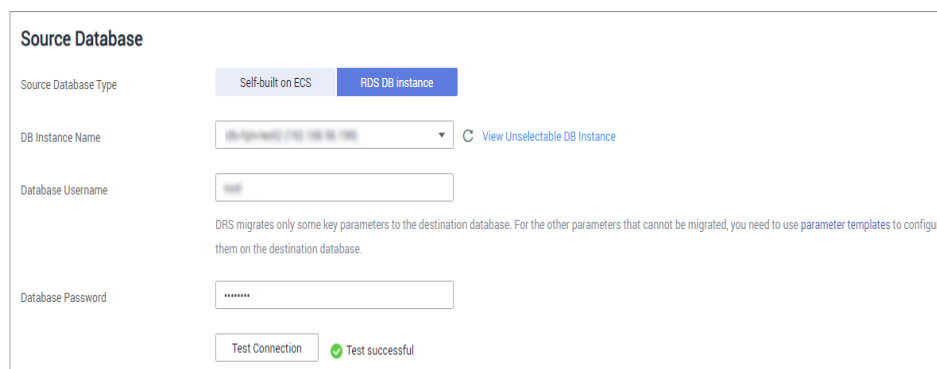| Parameter	Description
Source Database Type	Select **Self-built on ECS**.

Parameter	Description
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	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.
	
	NOTE
	  – The maximum size of a single certificate file that can be uploaded is 500 KB.
	  – If the SSL certificate is not used, your data may be at risk.

📖 **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 2-4 RDS DB instance - source database information

Table 2-6 RDS DB instance - source database information

Parameter	Description
Source 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.

- Destination database configuration

Figure 2-5 Destination database information

Table 2-7 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.
Database Password	The password for the database username.

Parameter	Description
Migrate Definer to User	– **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?** – **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**

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 2-6 Migration type

Table 2-8 Migration types and objects

Parameter	Description

Flow Control	You can choose whether to control the flow.
	
	● **Yes**
	You can customize the maximum migration speed.
	
	In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is **All day**. A maximum of three time ranges can be set, and they cannot overlap.
	
	The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.
	
	Figure 2-7 Flow control
	
	● **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.
	NOTE
	   – 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

Take Snapshot	If you perform a full migration, you can take a snapshot for your databases.
	
	● **No**
	This option applies to exports for which no data is written to the source database. If data is modified during a full migration, the exported data is point in time inconsistent. The stability and performance of a migration without a snapshot taken is better than that of a migration with a snapshot taken.
	
	● **Yes**
	A snapshot with consistent data at the point in time is generated during service running. Data changes during migration are not shown in the exported data.
	NOTE
	  – Snapshot reads use MySQL backup lock to lock global tables and automatically unlock them within 3s after consistent reads are enabled. To prevent full migration failures, take a snapshot when the source database is idle and does not perform DML or DDL operations during snapshot migration.
	  – Only MySQL full migration tasks support the snapshot mode. To use this function, you can submit a whitelist application.
	  – Do not perform DDL operations during migration in snapshot mode. Otherwise, full migration will fail.

Migrate Account	During a database migration, accounts need to be migrated separately. 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. ● **Yes** If you need to migrate accounts, see **Migrating Accounts**. ● **No** During migration, accounts, permissions, and passwords are not migrated.
Filter DROP DATABASE	To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations. ● 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
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **NOTE** ● 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. ● 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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 2-8 Pre-check

- 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 2-9 Modifying common parameters
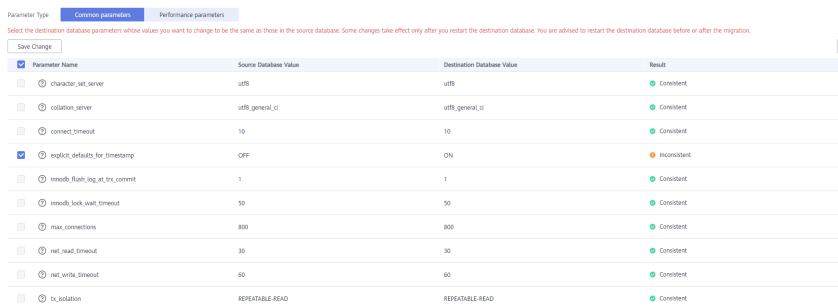
Performance parameter values in both the source and destination databases can be the same or different.

– 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:

i. Click **Use Source Database Value**.

DRS automatically makes the destination database values the same as those of the source database.

Figure 2-10 One-click modification

📖 **NOTE**

You can also manually enter parameter values.

ii. 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 2-11 Performance parameters

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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 2-12 Task startup settings

Table 2-9 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.

Parameter	Description
Synchronization Delay Threshold	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
	● 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 Notification**.
	● If the delay threshold is set to 0, no notifications will be sent to the recipient.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is **14**.
	NOTE
	Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, 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

2.2 From MySQL to DDM

Supported Source and Destination Databases

Table 2-10 Supported databases

Source DB	Destination DB
● On-premises MySQL databases ● MySQL databases on an ECS ● MySQL databases on other clouds ● RDS for MySQL	● DDM instances

Prerequisites

- You have logged in to the DRS console.

- Your account balance is greater than or equal to $0 USD.

- For details about the DB types and versions supported by real-time migration, see **Supported Databases**.

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see **Agency Management**.

Suggestions

> ⚠️ **CAUTION**
>
> - When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
>
> - 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.

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the migration to help you detect and resolve problems in advance.

- 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 50 MB/s during full migration, and two to four CPUs are occupied.

  - 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 certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-Level Comparison

  To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a migration task, read the following notes:

Table 2-11 Precautions

Type	Restrictions
Database permissions	- Full migration (minimum permissions): - The source database user must have the SELECT, SHOW VIEW, and EVENT permissions. - 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. - Full+incremental migration (minimum permissions): - The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. - The DDM destination database user must have the following permissions: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT. In addition, grant the SELECT permission on all tables. - The DDM destination database user must have the permission on the database to be migrated.
Source database	- The names of the source databases and tables cannot contain non-ASCII characters, or special characters <'>.`/\" - The source database name cannot be **ib_logfile**. - The binlog of the MySQL source database must be enabled and use the row-based format. - If the storage space is sufficient during incremental migration, store the source database binlog for as long as possible. The recommended retention period is three days. - If the **expire_logs_days** value of the source database is **0**, the migration may fail. - During an incremental migration, the **server-id** value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the **server_id** value ranges from **2** to **4294967296**. If the source database is MySQL 5.7 or later, the **server_id** value ranges from **1** to **4294967296**. - Enable **skip-name-resolve** for the MySQL source database to reduce the possibility of connection timeout. - Enable the Global Transaction Identifier (GTID) of the source database.

Type	Restrictions
Destination database	- Ensure that the destination database is empty before starting the migration. Otherwise, data with the same primary key in the destination may be overwritten during incremental migration. - The destination DB instance and associated RDS DB instance must be available. If the RDS DB instance type is primary/standby, the replication status must be normal. - The associated RDS DB instance must have sufficient storage space. - The character set of the associated RDS database must be the same as that of the source database. - If the destination 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. - The **AUTO_INCREMENT** value of a table in the destination cannot be less than that of a table in the source.
Migration objects	- Only the source database data can be migrated. The table structure and other database objects of the source database cannot be migrated. - Create table structures and indexes in the destination database that corresponds to the schema of the source database. Objects that are not created in the destination database are not to be migrated. - The table structure created in the destination database must be the same as that in the source database. - Tables with storage engine different to MyISAM and InnoDB tables cannot be migrated. - Tables without primary keys cannot be migrated.

Type	Restrictions
Precautions	• If the data types are incompatible, the migration may fail. • Cascade operations cannot be performed on tables with foreign keys. • If the source DB instance is an RDS MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. • If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7 and 5.7.28), you need to submit an O&M application for testing the SSL connection. • 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. • Set the **expire_log_day** parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. • During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. • During an incremental migration, do not modify the table structure to be migrated in the source database. • During an incremental migration of table-level objects, you are not advised to rename the tables. • DDL operations are not supported during the migration. Otherwise, the migration task may fail. • During an incremental migration, do not perform the restoration operation on the source database. • If the target DDM version is later than 3.0.4.1, DRS automatically updates the start value of the DDM sequence when the task is complete.

Procedure

Step 1 On the **Create Replication Instance** page, select a region, configure task details, description, and the replication instance, and click **Next**.

Figure 2-13 Migration task information

Table 2-12 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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 2-14 Replication instance information

Table 2-13 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**. ● VPC is suitable for migrations of cloud databases. ● VPN and Direct Connect are 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 Subnet** to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	● **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.
Enterprise Project	● If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. ● You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	● This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. ● After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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 configuration

Figure 2-15 Source database information

Table 2-14 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	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. **NOTE** – The maximum size of a single certificate file that can be uploaded is 500 KB. – If the SSL certificate is not used, your data may be at risk.

☐ **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 2-16 Destination database information

Table 2-15 Destination database settings

Parameter	Description
DB Instance Name	The DDM instance selected when you create the replication instance.
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 databases are encrypted and stored in DRS, and will be cleared after the task is deleted.

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

Figure 2-17 Migration object
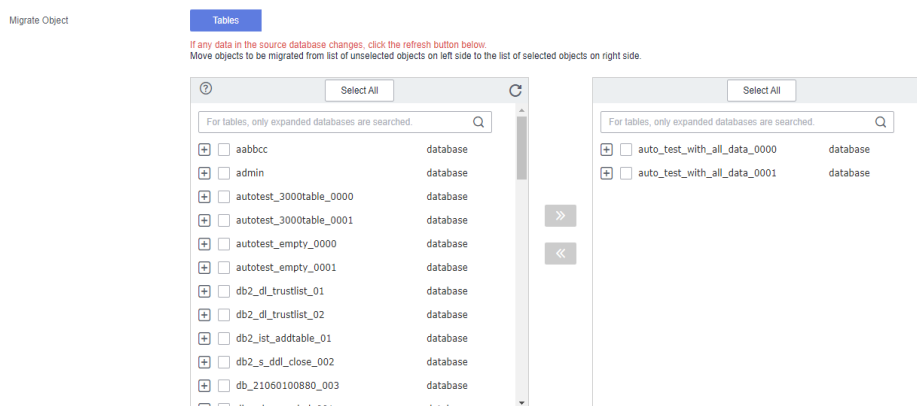
Table 2-16 Migration object

Parameter	Description
Migrate Object	You can migrate table-level objects to destination databases based on service requirements. 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. **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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 2-18 Task Check
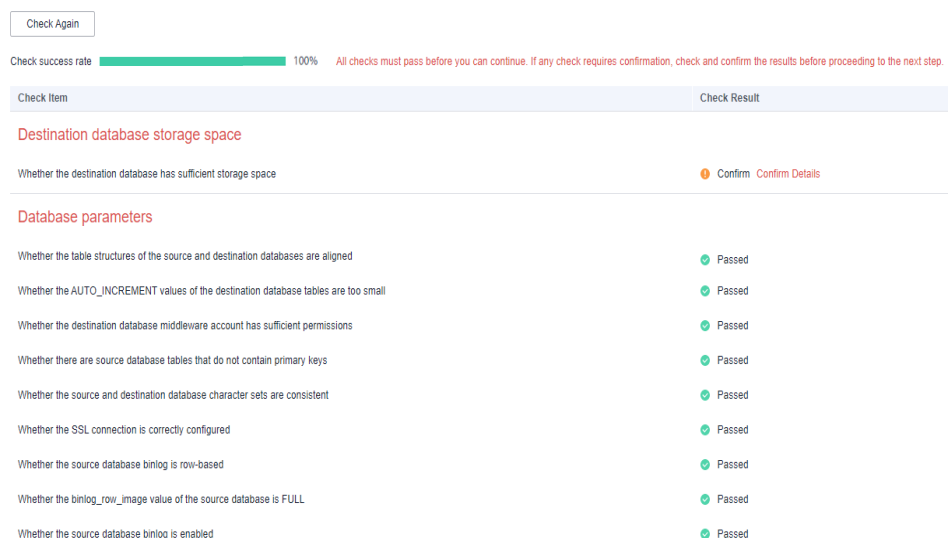
● If the check is complete and the check success rate is 100%, click **Next**.

📖 **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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 2-19 Task startup settings

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

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

* SMN Topic [ ] ▼ C ?

Synchronization Delay Threshold(s) ⊙ ?

* Stop Abnormal Tasks After [14] ? Abnormal tasks run longer than the period you set (unit: day) will automatically stop.
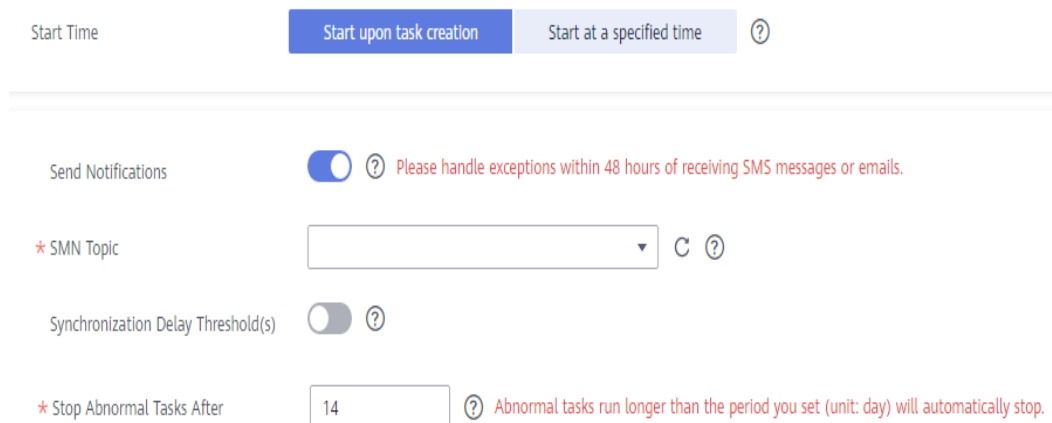
Table 2-17 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.
Synchronizat ion Delay Threshold	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** ● 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 Notification**. ● If the delay threshold is set to 0, no notifications will be sent to the recipient.

Parameter	Description
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is **14**. **NOTE** Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 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

2.3 From MySQL to GaussDB(for MySQL) Primary/Standby

Supported Source and Destination Databases

Table 2-18 Supported databases

Source DB	Destination DB
• On-premises MySQL databases • MySQL databases on an ECS • MySQL databases on other clouds • RDS for MySQL	• GaussDB(for MySQL) primary/standby

Prerequisites

- You have logged in to the DRS console.

- Your account balance is greater than or equal to $0 USD.

- For details about the DB types and versions supported by real-time migration, see **Supported Databases**.

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see **Agency Management**.

Suggestions

> **⚠ CAUTION**
>
> - When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
> - 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.

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the migration to help you detect and resolve problems in advance.

- 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 50 MB/s during full migration, and two to four CPUs are occupied.

  - 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 certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-Level Comparison

  To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a migration task, read the following notes:

Table 2-19 Precautions

Type	Restrictions

Database permissions	- Full migration (minimum permissions):
	   – The source database user must have the SELECT, SHOW VIEW, and EVENT permissions.
	   – The destination database user must have the following permissions: SELECT, CREATE, 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.
	 - Full+incremental migration (minimum permissions):
	   – The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT.
	   – The destination database user must have the following permissions: SELECT, CREATE, 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.

Migration objects	- Supported objects: databases, tables, views, indexes, constraints, functions, stored procedures, triggers, and events. - The system database and event statuses cannot be migrated. - Encrypted tables cannot be migrated. - Tables with storage engine different to MyISAM and InnoDB tables cannot be migrated.

Type	Restrictions
Source database	- The database name, table name, and view name in the source database cannot contain non-ASCII characters or the following characters .<'>`/\" - The binlog of the MySQL source database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. - If the **expire_logs_days** value of the source database is **0**, the migration may fail. - During an incremental migration, the **server_id** value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the **server_id** value ranges from **2** to **4294967296**. If the source database is MySQL 5.7 or later, the **server_id** value ranges from **1** to **4294967296**. - Enable **skip-name-resolve** for the MySQL source database to reduce the possibility of connection timeout. - Enable the Global Transaction Identifier (GTID) of the source database. - The source cannot contain empty databases.
Destination database	- Data cannot be migrated from a newer version database to an older version database. - You are advised to use the row-based binlog in the destination MySQL database. Otherwise, an error may occur during an incremental migration. - The destination DB instance is running properly. - The destination DB instance must have sufficient storage space. - The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database). - The destination database isolation level must be set to at least read committed.

Type	Restrictions
Precautions	- Cascade operations cannot be performed on tables with foreign keys. - If the source databases are RDS MySQL instances, TDE is not supported, and tables with the encryption function cannot be created. - 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 and 5.7.28), you need to submit an O&M application for testing the SSL connection. - 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 destination database of a migration task can be set to **Read-only** or **Read/Write**. **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 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 **expire_log_day** parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. - During task startup or full migration, you are not advised to perform DDL operations on the source database, such as deleting the database, indexes, or views. Otherwise, the migration may fail. - During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. - To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL and DML operations) during migration. - Ensure that the source database is publicly accessible before the migration task is complete. - During migration, do not write the statement-based binlog into the source database. - During migration, do not clear the binlog in the source database.

Type	Restrictions
	- During migration, GaussDB(for MySQL) automatically converts the MyISAM table to the InnoDB table. If the conversion fails, the migration fails. - During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database. - During an incremental migration, if distributed transactions exist in the source database, the migration may fail. - 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 of table-level objects, you are not advised to rename the tables. - When a migration task is completed, the system begins to migrate the selected events and triggers. You must therefore check the status of the migration log to ensure database integrity.

Procedure

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

Step 2 On the **Create Replication Instance** page, select a region, configure task details, description, and the replication instance, and click **Next**.

Figure 2-20 Migration task information

Table 2-20 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.

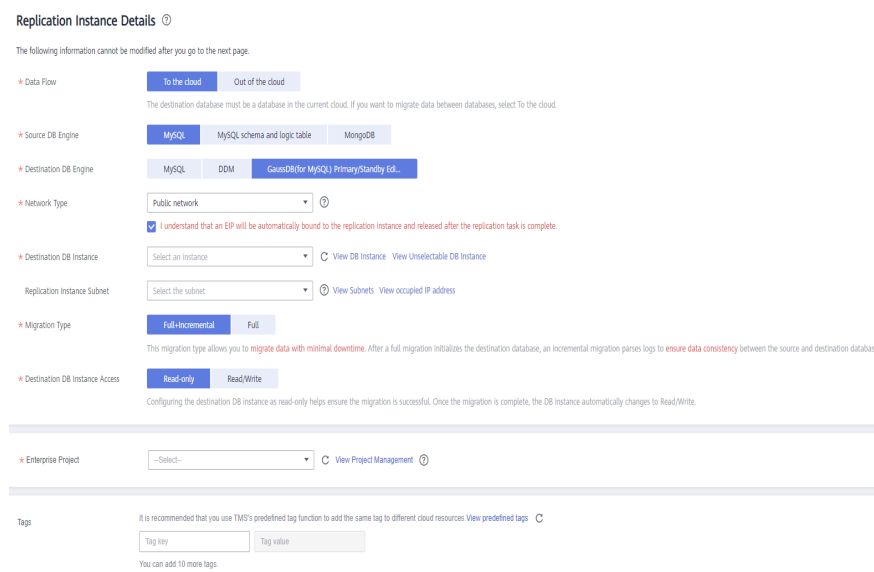| Parameter	Description
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 2-21 Replication instance information

Table 2-21 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) Primary/Standby Edition**.
Network Type	Available options: **VPC**, **Public network**, and **VPN or Direct Connect**. By default, the value is **Public network**. ● VPC is suitable for migrations of cloud databases. ● VPN and Direct Connect are 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.

Parameter	Description
Destination DB Instance	The GaussDB(for MySQL) primary/standby instance you created.
Replication Instance Subnet	The subnet where the replication instance resides. You can also click **View Subnet** to go to the network console to view the subnet where the instance resides.
	By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Migration Type	● **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.

Destination Database Access	● **Read-only** During 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 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. The task cannot be modified after being created.

Parameter	Description
Enterprise Project	• If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. • You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	• This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. • After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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 2-22 Source database information

Table 2-22 Source database settings

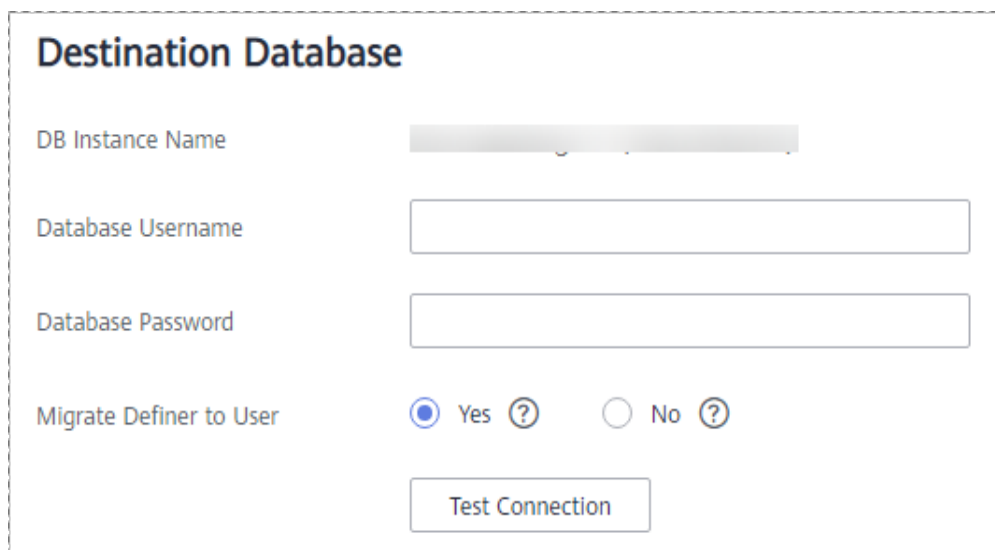| Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535

Parameter	Description
Database Username	The username for accessing the source database.
Database Password	The password for the database username.

SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.
	NOTE
	  – The maximum size of a single certificate file that can be uploaded is 500 KB.
	  – If the SSL certificate is not used, your data may be at risk.

📖 **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 2-23 Destination database information

Table 2-23 Destination database settings

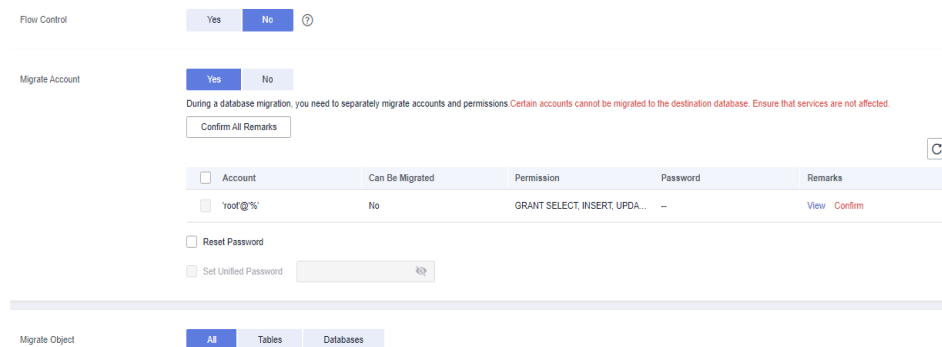| Parameter	Description
DB Instance Name	The GaussDB(for MySQL) primary/standby instance selected during the migration task creation and cannot be changed.
Database Username	The username for accessing the destination database.

Parameter	Description
Database Password	The password for the database username.
Migrate Definer to User	– **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?** **NOTE** For the migration from MySQL to GaussDB(for MySQL) primary/standby, you can only select **Yes**. After the migration, the Definers of all source database objects will be migrated to the user.

 NOTE

The username and password of the destination databases 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 2-24 Migration object

Table 2-24 Migration object

| Parameter	Description

Flow Control	You can choose whether to control the flow.
	● **Yes**
	You can customize the maximum migration speed.
	In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is **All day**. A maximum of three time ranges can be set, and they cannot overlap.
	The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.
	
	Figure 2-25 Flow control
	
	● **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.
	NOTE
	  – 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 Account	During a database migration, accounts need to be migrated separately. 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. ● **Yes** If you need to migrate accounts, see **Migrating Accounts**. ● **No** During migration, accounts, permissions, and passwords are not migrated.

Parameter	Description
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **NOTE** ● 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. ● 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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 2-26 Pre-check
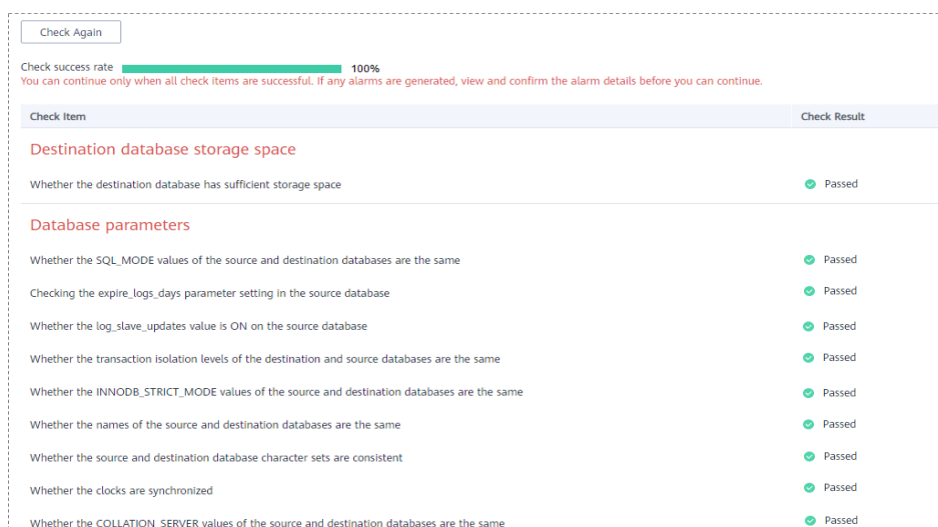
- If the check is complete and the check success rate is 100%, click **Next**.

  ☐ **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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 2-27 Task startup settings
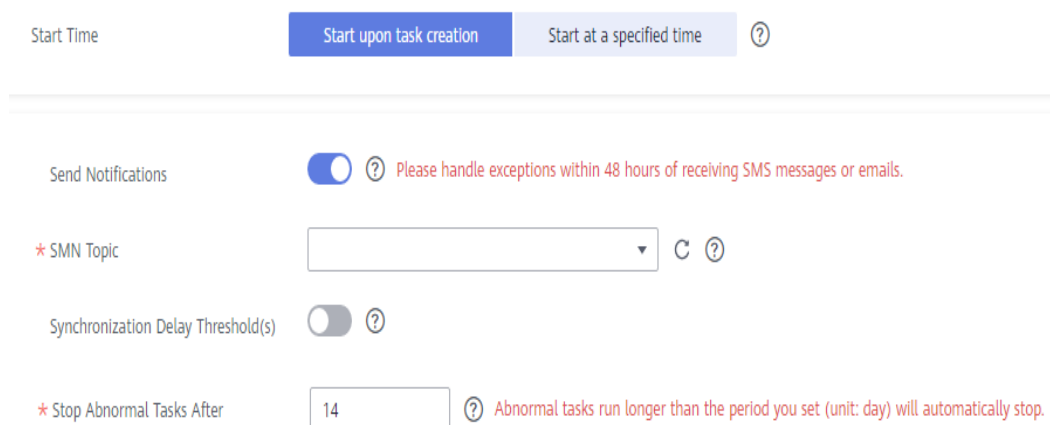
Table 2-25 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.

Parameter	Description
Synchronization Delay Threshold	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** ● 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 Notification**. ● If the delay threshold is set to 0, no notifications will be sent to the recipient.
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.

----End

2.4 From MySQL Schema and Logic Table to DDM

Supported Source and Destination Databases

Table 2-26 Supported databases

Source DB	Destination DB
● DDM instances ● On-premises MyCAT middleware ● MyCAT middleware on an ECS	● DDM instances

Prerequisites

- You have logged in to the DRS console.

- Your account balance is greater than or equal to $0 USD.

- For details about the DB types and versions supported by real-time migration, see **Supported Databases**.

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see **Agency Management**.

Suggestions

> ⚠ **CAUTION**
>
> - When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
>
> - 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.

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the migration to help you detect and resolve problems in advance.

- 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 50 MB/s during full migration, and two to four CPUs are occupied.

  - 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 certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-Level Comparison

  To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a migration task, read the following notes:

Table 2-27 Precautions

Type	Restrictions
Database permissions	- Full migration (minimum permissions): - The source sharded database user must have the SELECT, SHOW VIEW, and EVENT permissions. - 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. - Full+incremental migration (minimum permissions): - The source sharded database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT. - The DDM destination database user must have the following permissions: CREATE, DROP, ALTER, INDEX, INSERT, DELETE, UPDATE, and SELECT. In addition, grant the select permission on all tables. - The DDM destination database user must have the permission on the database to be migrated.
Source database	- The database names and table names of the source sharding middleware cannot contain the following characters: '<>/\ and non-ASCII characters. - The binlog of the MySQL source database must be enabled and use the row-based format. - If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. - If the **expire_logs_days** value of the source database is **0**, the migration may fail. - During an incremental migration, the **server-id** value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the **server_id** value ranges from **2** to **4294967296**. If the source database is MySQL 5.7 or later, the **server_id** value ranges from **1** to **4294967296**. - Enable **skip-name-resolve** for the MySQL source database to reduce the possibility of connection timeout. - Enable the Global Transaction Identifier (GTID) of the source database.

Type	Restrictions
Destination database	- Ensure that the destination database is empty before starting the migration. Otherwise, data in the destination may be overwritten during incremental migration. - The destination DB instance and associated RDS DB instance must be available. If the RDS DB instance type is primary/standby, the replication status must be normal. - The associated RDS DB instance must have sufficient storage space. - The character set of the associated RDS database must be the same as that of the source database. - If the destination 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. - The **AUTO_INCREMENT** value of a table in the destination cannot be less than that of a table in the source.
Migration objects	- Only the source database data can be migrated. The table structure and other database objects of the source database cannot be migrated. - Create table structures and indexes in the destination database that corresponds to the schema of the source database. Objects that are not created in the destination database are not to be migrated. - The table structure created in the destination database must be the same as that in the source database. - If the source database is a DDM database, the table cannot contain sharding keys of the timestamp type. - Tables with storage engine different to MyISAM and InnoDB tables cannot be migrated. - Tables without primary keys cannot be migrated.

Type	Restrictions
Precautions	• 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 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.
	• Set the **expire_log_day** parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.
	• During the task startup or full migration, you are not advised to perform DDL operations, such as deletion, on the source database. Otherwise, the migration may fail.
	• During the migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.
	• During an incremental migration, do not modify the table structure to be migrated 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.
	• During an incremental migration of table-level objects, you are not advised to rename the tables.
	• During an incremental migration, do not perform the restoration operation on the source database.
	• If the target DDM version is later than 3.0.4.1, DRS automatically updates the start value of the DDM sequence when the task is complete.

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, specify task name, description, and the replication instance details, and click **Next**.

Figure 2-28 Migration task information

Table 2-28 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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 2-29 Replication instance information

Table 2-29 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 **MySQL sharding**.
Destination DB Engine	Select **DDM**.
Network Type	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**. ● VPC is suitable for migrations of cloud databases. ● Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP. ● VPN is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. ● Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions.
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 Subnet** to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	• **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	The default minimum number of source DB instances is 2 and the maximum number is 16. You can set this parameter based on the number of source database shards.
Enterprise Project	• If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. • You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	• This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. • After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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 2-30 Source database information

Table 2-30 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	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. **NOTE** – The maximum size of a single certificate file that can be uploaded is 500 KB. – If the SSL certificate is not used, your data may be at risk.
Sharded Database	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 2-31 Destination database information

Table 2-31 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 databases 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 2-32 Migration object
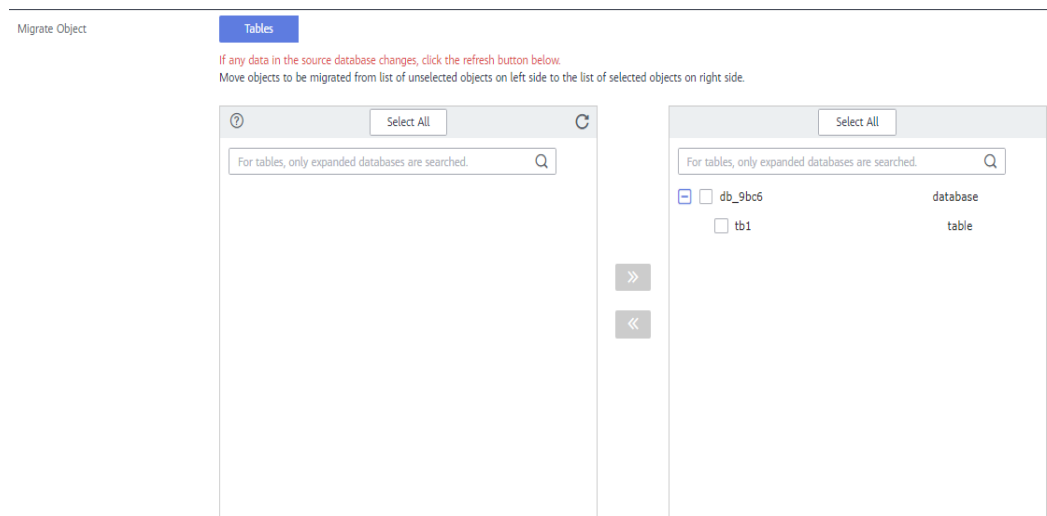
Table 2-32 Migrate Object

Parameter	Description
Migrate Object	After the objects are migrated to the destination database, the object names will remain the same as those in the source database and cannot be modified.
	NOTE Structures that are not created in the destination database are not migrated by default.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 2-33 Pre-check
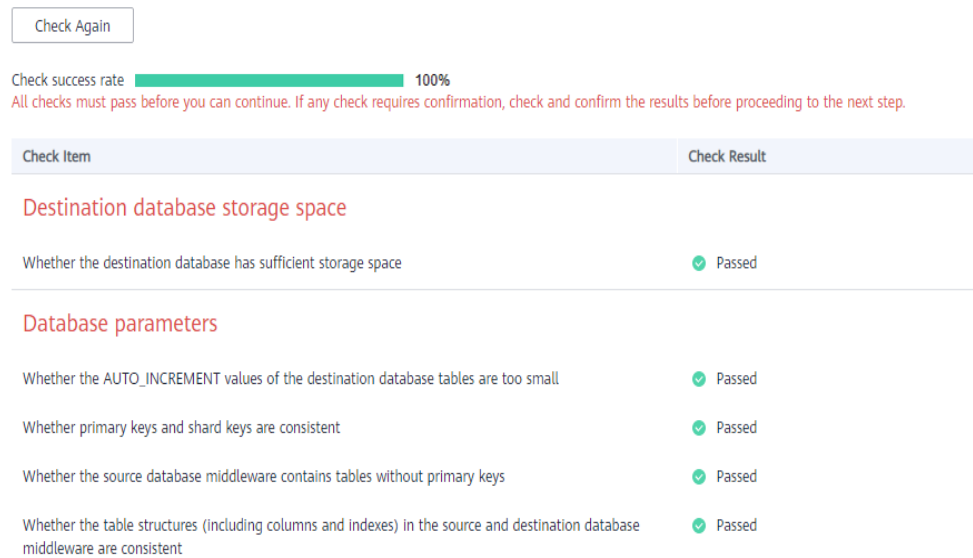
- If the check is complete and the check success rate is 100%, click **Next**.

  ☐ **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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

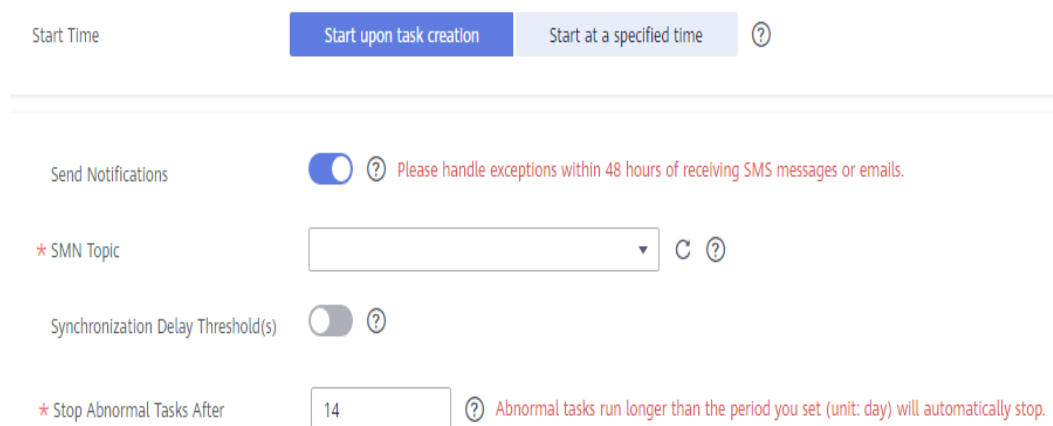**Figure 2-34** Task startup settings

Table 2-33 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.
Synchronization Delay Threshold	During an incremental 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** ● 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 Notification**. ● If the delay threshold is set to 0, no notifications will be sent to the recipient.
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.

----End

2.5 From MongoDB to DDS

Supported Source and Destination Databases

Table 2-34 Supported databases

Source DB	Destination DB
● On-premises MongoDB databases ● MongoDB databases on an ECS ● MongoDB database on other clouds ● DDS instances	● DDS instances

Prerequisites

- You have logged in to the DRS console.

- Your account balance is greater than or equal to $0 USD.

- For details about the DB types and versions supported by real-time migration, see **Supported Databases**.

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see **Agency Management**.

Suggestions

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the 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.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-Level Comparison

  To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a migration task, read the following notes:

Table 2-35 Precautions

Type	Restrictions
Database permissions	Source database (minimum permissions): ● Full migration: – 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: – 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. Minimum permission requirements: The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database. If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.

Type	Restrictions
Migration objects	• 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.
Source database	• The source database name cannot contain /\."$ or spaces. The collection name or view name cannot start with **system.** or contain the dollar sign ($). • If the incremental source data cluster is migrated, the source database balancer must be disabled. • The source cannot be a GaussDB(for Mongo) instance.
Destination database	• The destination DB instance is running properly. • The destination DB instance must have sufficient storage space. • When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique. • 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. • Data cannot be migrated from a newer version database to an older version database.

Type	Restrictions
Precautions	- Objects that have dependencies must be migrated at the same time to avoid migration failure. Common dependencies: collections referenced by views, and views referenced by views - Replica set: The MongoDB replica set instance must be available and have primary nodes. - Source database from a single node instance on other clouds cannot be migrated. - If the source database is not on a cluster instance, the following operations and commands are supported during incremental migration: - Creating and deleting databases - Adding, deleting, and updating documents - Creating and deleting collections - Creating and deleting indexes - Creating and deleting views - The convertToCapped, collMod, and renameCollection 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 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: - 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 DBPointer and DBRef 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 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. - The value of **block_compressor** is determined by **stats().wiredTiger.creationString.block_compressor** of the

Type	Restrictions
	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 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.
	● Currently, DCC does not support DDS DB instances. Migration tasks cannot be created.
	● 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, enter the information about all primary and secondary shard nodes in the same cluster to reduce the impact of a primary/secondary switchover on the migration task.
	● 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. – 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.

Type	Restrictions
	• To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL and DML operations) during migration.
	• During the migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.
	• During task startup or full migration, you are not advised to perform DDL operations on the source database, such as deleting databases, collections, indexes, documents, or views. Otherwise, the migration may fail.
	• During migration, data rollback caused by a primary/standby switchover of the source database is not supported.
	• During an incremental migration of collections, you are advised not to rename the collections.
	• Files larger than 16 MB cannot be inserted to or updated in the source database during full or incremental migration.

	• In the incremental migration phase, concurrent replay is performed at the collection level to maintain the migration performance. In the following scenarios, only single-thread write is supported and concurrent replay is not supported:
	  – The collection index contains a unique key.
	  – The value of **capped** of the collection attribute is **true**.
	In either of the preceding scenarios, the task delay may increase.

	• 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.

Procedure

This section uses MongoDB sharded clusters as an example to describe how to configure a task for migrating MongoDB databases 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, specify the task name, description, and the replication instance details, and click **Next**.

Figure 2-35 Migration task information

Table 2-36 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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 2-36 Replication instance information

Table 2-37 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**. ● VPC is suitable for migrations of cloud databases. ● Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP. ● VPN is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. ● Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions.
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 Subnet** to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Migration Type	● **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.
Source DB Instance Type	If you select **Full+Incremental** for **Migration Type**, set this parameter based on the source database. ● 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**.
Obtain Incremental Data	This parameter is available for configuration if **Source DB Instance Type** is set to **Cluster**. You can determine how to capture data changes during the incremental synchronization. ● 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 must disable the balancer of the source database. For details, see **How Do I Disable the Balancer?** When testing the connectivity between the source and the DRS instance, you need to enter the connection information of each shard in the source database on the task configuration page. ● 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. **NOTE** Currently, only whitelisted users can use this function.

Parameter	Description
Source Shard Quantity	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. The number of source shards ranges from 2 to 32. Specify this parameter based on the actual number of shards in the source DB.
Enterprise Project	• If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. • You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	• This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. • After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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

Table 2-38 Source database information

Parameter	Description
mongos Address	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 - 65534 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. **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.
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is **admin**.
mongos Username	The username for accessing the source database.
mongos Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. **NOTE** – The maximum size of a single certificate file that can be uploaded is 500 KB. – If the SSL certificate is not used, your data may be at risk.
Sharded Database	Enter the information about the sharded databases in the source database. If the source is a DDS instance, apply for an IP address for the shard node by referring to **Document Database Service User Guide**.

☐ **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

Table 2-39 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 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 4 On the **Set Task** page, select migration objects and click **Next**.
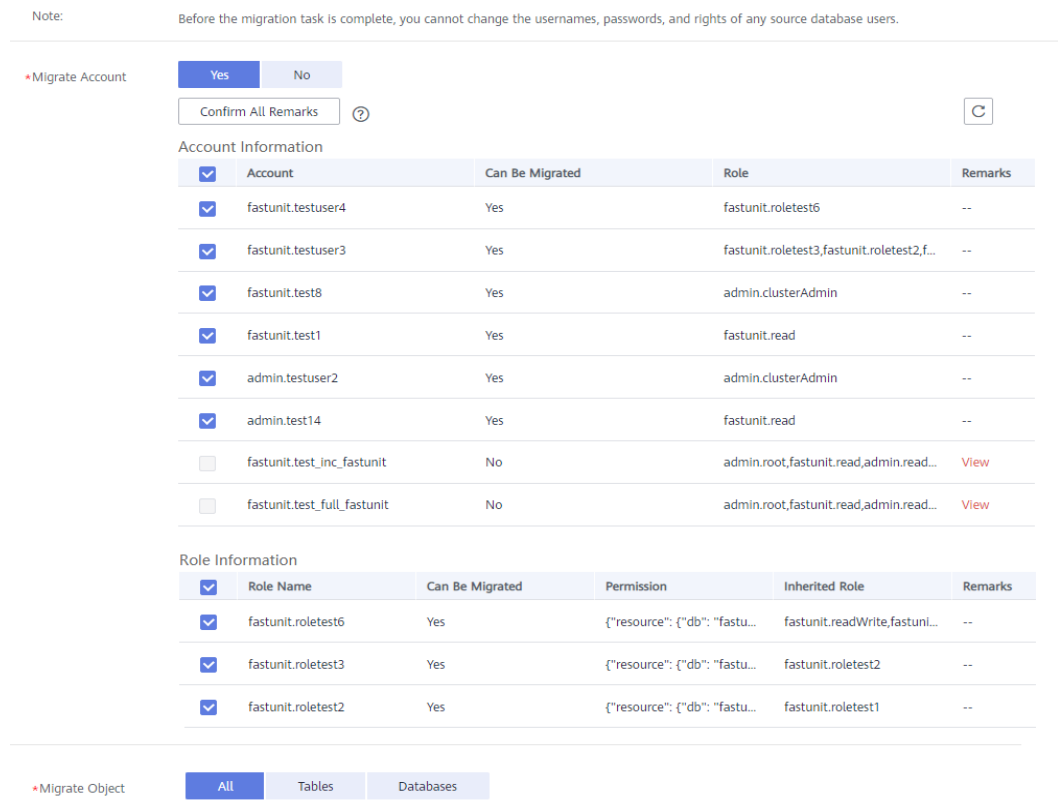
Figure 2-39 Migration object

Table 2-40 Migrate Object

| Parameter	Description
Migrate Account	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. ● **Yes** If you need to migrate accounts, see **Migrating Accounts**. ● **No** During the migration, accounts and roles are not migrated.

Parameter	Description
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **NOTE** ● 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. ● 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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 2-40 Pre-check
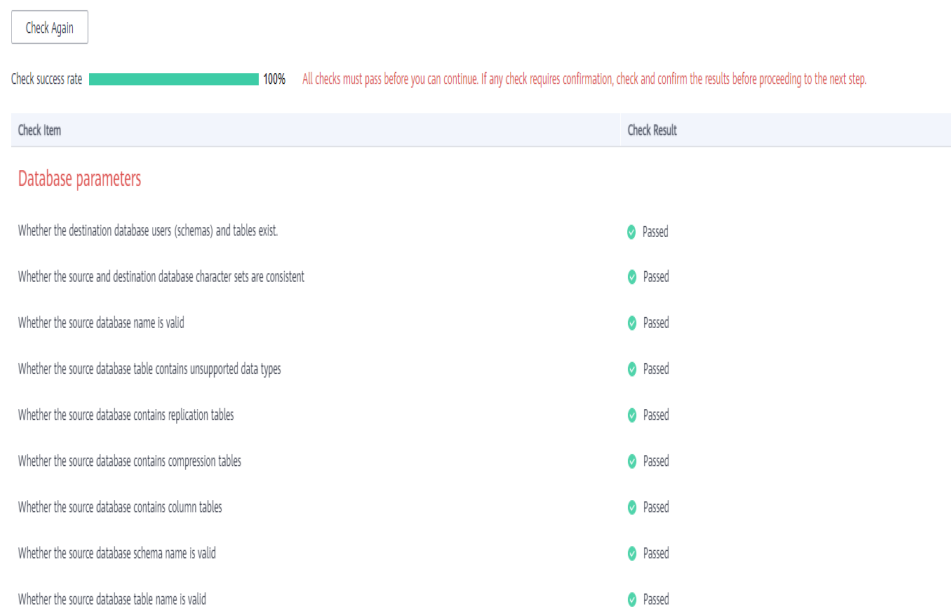
- If the check is complete and the check success rate is 100%, click **Next**.

  ☐ **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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

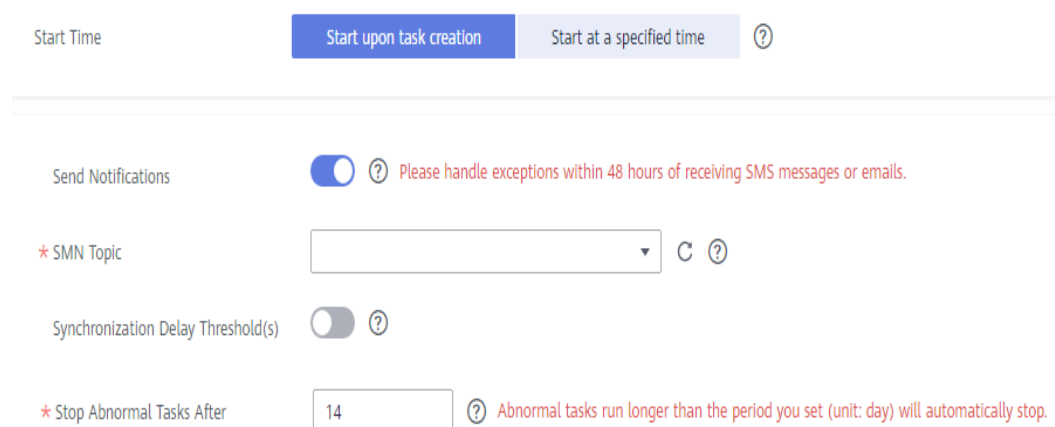**Figure 2-41** Task startup settings

Table 2-41 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.
Synchronization Delay Threshold	During an incremental 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** ● 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 Notification**. ● If the delay threshold is set to 0, no notifications will be sent to the recipient.
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.

----End

2.6 From MongoDB to GaussDB(for Mongo)

Supported Source and Destination Databases

Table 2-42 Supported Databases

Source DB	Destination DB
- On-premises MongoDB databases - MongoDB databases on an ECS - MongoDB database on other clouds	GaussDB(for Mongo) instances

Prerequisites

- You have logged in to the DRS console.

- Your account balance is greater than or equal to $0 USD.

- For details about the DB types and versions supported by real-time migration, see **Supported Databases**.

- If a subaccount is used to create a DRS task, ensure that an agency has been added. To create an agency, see **Agency Management**.

Suggestions

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the 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.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-Level Comparison

  To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a migration task, read the following notes:

Table 2-43 Precautions

Type	Restrictions

Database permissions	Source database (minimum permissions):
	● Full migration:
	  – 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:
	  – 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.
	Minimum permission requirements: The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database. If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.

Type	Restrictions
Migration objects	- 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 **_id** field without indexes are not supported. - The first parameter of **BinData()** cannot be **2**.
Source database	- The source database name cannot contain /\."$ or spaces. The collection name and view name cannot start with **system.** or contain the dollar sign ($). - If the incremental source data cluster is migrated, the source database balancer must be disabled. - The source cannot be a GaussDB(for Mongo) instance.
Destination database	- The destination DB instance is running properly. - The destination DB instance must have sufficient storage space. - When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique. - 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. - Data cannot be migrated from a newer version database to an older version database.

Type	Restrictions
Precautions	- Objects that have dependencies must be migrated at the same time to avoid migration failure. Common dependencies: collections referenced by views, and views referenced by views - Replica set: The MongoDB replica set instance must be available and have primary nodes. - If the source is not a cluster instance, you can perform the following operations during incremental migration: - Creating and deleting databases - Adding, deleting, and updating documents - Creating and deleting collections - Creating and deleting indexes - Creating and deleting views - The convertToCapped, collMod, and renameCollection commands are supported. - During a full plus incremental migration from a replica set to a cluster or between clusters, the objects to be migrated cannot be deleted. Otherwise, the migration task will fail. - 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 **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, enter the information about all primary and secondary shard nodes in the same cluster to reduce the impact of a primary/secondary switchover on the migration task. - The maximum number of collections that can be created in a GaussDB(for Mongo) cluster is calculated as follows: Maximum number of chunks x 4 x Number of shards

Type	Restrictions
	The following table lists the maximum number of chunks of different GaussDB(for Mongo) instance specifications.
	1 vCPUs, 4 GB -> 50
	2 vCPUs, 8 GB -> 100
	4 vCPUs, 16 GB -> 200
	8 vCPUs, 32 GB -> 400
	16 vCPUs, 64 GB -> 800
	32 vCPUs, 128 GB -> 1600
	Check whether the destination database meets the requirements based on the number of source database collections.
	● 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.
	– 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 ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL and DML operations) during migration.
	● During the migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.
	● During task startup or full migration, you are not advised to perform DDL operations on the source database, such as deleting databases, collections, indexes, documents, or views. Otherwise, the migration may fail.
	● During migration, data rollback caused by a primary/standby switchover of the source database is not supported.
	● During an incremental migration of collections, you are advised not to rename the collections.
	● Files larger than 16 MB cannot be inserted to or updated in the source database during full or incremental migration.
	● 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.

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
• RDS for MySQL	• On-premises MySQL databases • MySQL databases on an ECS • MySQL databases on other clouds • RDS for MySQL

Prerequisites

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

Suggestions

> ⚠️ **CAUTION**
>
> - When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
> - 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.

- 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.

- 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 you have to migrate data during peak hours, you can select **Yes** for **Flow Control** to adjust the migration speed.

  - 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.

  - 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%.

  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.

  - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.

  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**

- Data-level comparison

  To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

Precautions

Before creating a migration task, read the following notes:

Table 3-2 Precautions

Type	Restrictions
Precautions	• Data cannot be migrated from a newer version database to an older version database.
	• If the DCC does not support instances with 4 vCPUs and 8 GB memory or higher instance specifications, the migration task cannot be created.
	• Objects that have dependencies must be migrated at the same time to avoid migration failure. Common dependencies: tables referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys
	• Cascade operations cannot be performed on tables with foreign keys.
	• The primary/standby switchover is not supported if GTID is not enabled for the source database.
	• If the source and destination sides are RDS MySQL instances, transparent data encryption (TDE) is not supported, and tables with the encryption function cannot be created.
	• 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 and 5.7.28), you need to submit an O&M application for testing the SSL connection.
	• Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key.
	• During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.
	• During the migration, you are not advised to perform a primary/standby switchover on the destination database.
	• Ensure that the destination database is publicly accessible before the migration task is complete.
	• 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 full migration, DDL operations such as table structure modification are not supported.
	• The destination DB instance must be readable and writable. Otherwise, the migration may fail.

Procedure

This section uses the migration from an RDS 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, configure task details, description, and the replication instance, and click **Next**.

Figure 3-1 Migration task information

Table 3-3 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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 3-2 Replication instance information

Table 3-4 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** ● VPC is suitable for migrations of cloud databases. ● VPN and Direct Connect are 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 Subnet** to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description

Migration Type	● **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.

Enterprise Project	● If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. ● You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	● This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. ● After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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 3-3 Source database information

Table 3-5 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 **Migration Information** area on the **Basic Information** page, click **Update Password** next to the **Source Database Password** field. In the displayed dialog box, change the password. This action only updates DRS with the changed password.

📖 **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 3-4 Destination database information

Table 3-6 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 **Migration Information** area on the **Basic Information** page, click **Update Password** next to the **Destination Database Password** field. In the displayed dialog box, change the password. This action only updates DRS with the changed password.

Parameter	Description
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. **NOTE** ● The maximum size of a single certificate file that can be uploaded is 500 KB. ● If the SSL certificate is not used, your data may be at risk.
Migrate Definer to User	● **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?** ● **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**

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 3-5 Migration Type

Table 3-7 Migration types and objects

Parameter	Description
Flow Control	You can choose whether to control the flow. ● **Yes** You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is **All day**. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s. **Figure 3-6** Flow control ● **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. **NOTE** – 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

Take Snapshot	If you perform a full migration, you can take a snapshot for your databases.
	
	● **No**
	This option applies to exports for which no data is written to the source database. If data is modified during a full migration, the exported data is point in time inconsistent. The stability and performance of a migration without a snapshot taken is better than that of a migration with a snapshot taken.
	
	● **Yes**
	A snapshot with consistent data at the point in time is generated during service running. Data changes during migration are not shown in the exported data.
	NOTE
	  – Snapshot reads use MySQL backup lock to lock global tables and automatically unlock them within 3s after consistent reads are enabled. To prevent full migration failures, take a snapshot when the source database is idle and does not perform DML or DDL operations during snapshot migration.
	  – The snapshot function is supported only for full migration of databases whose engine is MySQL.

Filter DROP DATABASE	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. The database deletion operation can be filtered by default. ● 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. **NOTE** Currently, only the full plus incremental migrations from RDS MySQL to MySQL are supported.
Migrate Account	During a database migration, accounts need to be migrated separately. 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. ● **Yes** If you need to migrate accounts, see **Migrating Accounts**. ● **No** During the migration, accounts, permissions, and passwords are not migrated.

Parameter	Description
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **NOTE** ● 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. ● 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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 3-7 Pre-check
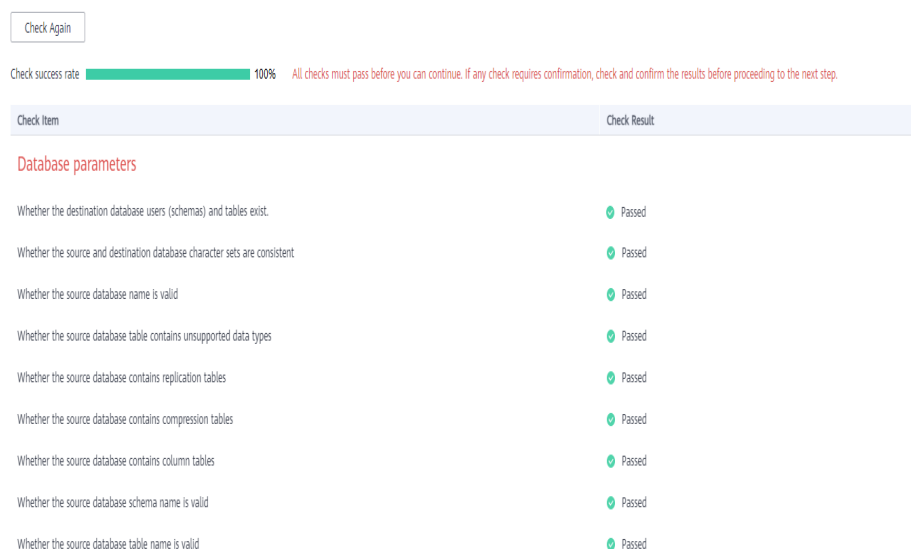
- If the check is complete and the check success rate is 100%, click **Next**.

  ☐ 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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 3-8 Task startup settings

Table 3-8 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.

Parameter	Description
Synchronization Delay Threshold	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** • 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 Notification**. • If the delay threshold is set to 0, no notifications will be sent to the recipient.
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.

----End

3.2 From DDS to MongoDB

Supported Source and Destination Databases

Table 3-9 Supported Databases

Source DB	Destination DB
• DDS instances	• On-premises MongoDB databases • MongoDB databases on an ECS • MongoDB database on other clouds

Prerequisites

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

Suggestions

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the 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.
  - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
  - For more information about the impact of DRS on databases, see **What Is the Impact of DRS on Source and Destination Databases?**
- Data-Level Comparison

  To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a migration task, read the following notes:

Table 3-10 Precautions

Type	Restrictions
Database permissions	Source database (minimum permissions): ● Full migration: – 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: – 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. Minimum permission requirements: The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database. If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database and the read permission for the config database.

Type	Restrictions
Precautions	• Data cannot be migrated from a newer version database to an older version database.
	• System databases cannot be migrated. The username and role must be manually created in the destination database.
	• 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.
	• 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 storage engine of the destination database is not WiredTiger, DRS does not migrate compression parameters.
	• Collections that contain the **_id** field without indexes are not supported.
	• The first parameter of **BinData()** cannot be **2**.
	• The statement for creating a view cannot contain a regular expression.
	• 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.
	• During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.
	• To ensure data consistency, you are not allowed to modify the destination database (including but not limited to DDL and DML operations) during the entire migration process.
	• During migration, data rollback caused by a primary/standby switchover of the source database is not supported.

Type	Restrictions

	• Files larger than 16 MB cannot be inserted to or updated in the source database during full or incremental migration.
	
	• In the incremental migration phase, concurrent replay is performed at the collection level to maintain the migration performance. In the following scenarios, only single-thread write is supported and concurrent replay is not supported:
	  – The collection index contains a unique key.
	  – The value of **capped** of the collection attribute is **true**.
	In either of the preceding scenarios, the task delay may increase.
	
	• 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 describes how to migrate 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, specify the task name, description, and the replication instance details, and click **Next**.

Figure 3-9 Migration task information

Table 3-11 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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).

Parameter	Description
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 3-10 Replication instance information

Table 3-12 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** • VPC is suitable for migrations of cloud databases. • VPN and Direct Connect are 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.

Parameter	Description
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 Subnet** to go to the network console to view the subnet where the instance resides.
	By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.
Migration Type	• **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. **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.
Enterprise Project	• If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. • You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	• This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. • After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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 3-11 Source database information

Table 3-13 Source database settings

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

 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 3-12 Destination database information

Table 3-14 Destination 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. The port of the source database. Range: 1 - 65534
	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 instance.
	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.
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.

Parameter	Description
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. **NOTE** ● The maximum size of a single certificate file that can be uploaded is 500 KB. ● If the SSL certificate is not used, your data may be at risk.

📖 **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 3-13 Migration object
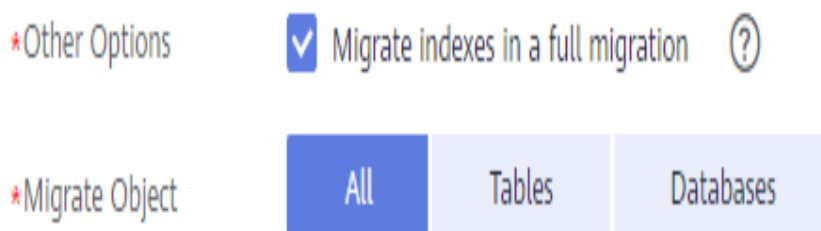
Table 3-15 Migrate Object

| Parameter	Description
Other Options	Determine whether to migrate the indexes you create during full migration. The default index based on _id is automatically created in the destination. If indexes are not migrated, the indexes are not compared.

Parameter	Description
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 3-14 Pre-check
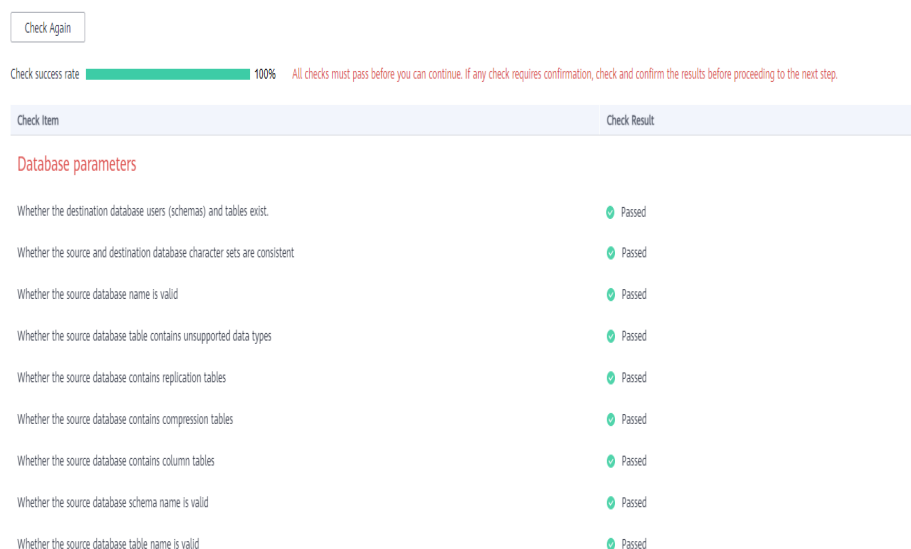
● If the check is complete and the check success rate is 100%, click **Next**.

📖 **NOTE**

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

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

Figure 3-15 Task startup settings

| Start Time | Start upon task creation | Start at a specified time | ❓ |

Send Notifications	🔵 ❓ Please handle exceptions within 48 hours of receiving SMS messages or emails.
✱ SMN Topic	[▼] C ❓
Synchronization Delay Threshold(s)	⚪ ❓
✱ Stop Abnormal Tasks After	[14] ❓ Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 3-16 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.

Parameter	Description
Synchronization Delay Threshold	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** ● 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 Notification**. ● If the delay threshold is set to 0, no notifications will be sent to the recipient.
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.

**----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
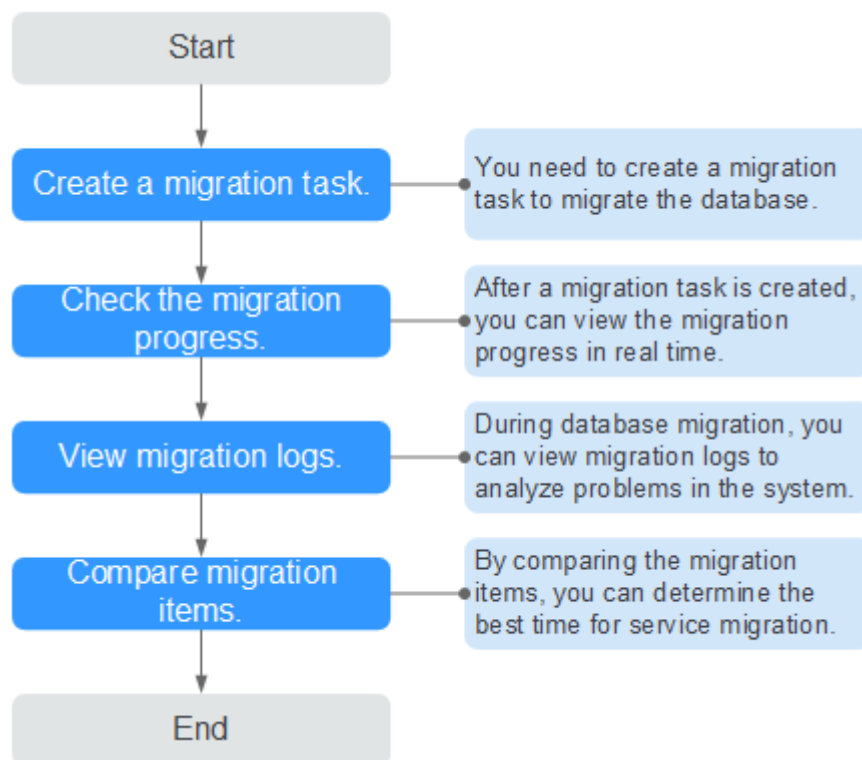
- **Step 1: Create a migration task.** Select the source and destination databases as required and create a migration task.

- **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 MySQL as an example to describe how to configure a migration task over a VPC network on the DRS console.

A VPC network is suitable for migrations of cloud databases.

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. 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, configure task details, description, and the replication instance, and click **Next**.

Figure 4-2 Migration task information

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.
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), digits, hyphens (-), and underscores (_).
Description	The description can contain up to 256 characters and cannot contain special characters !=<>&'\"

Figure 4-3 Replication instance information
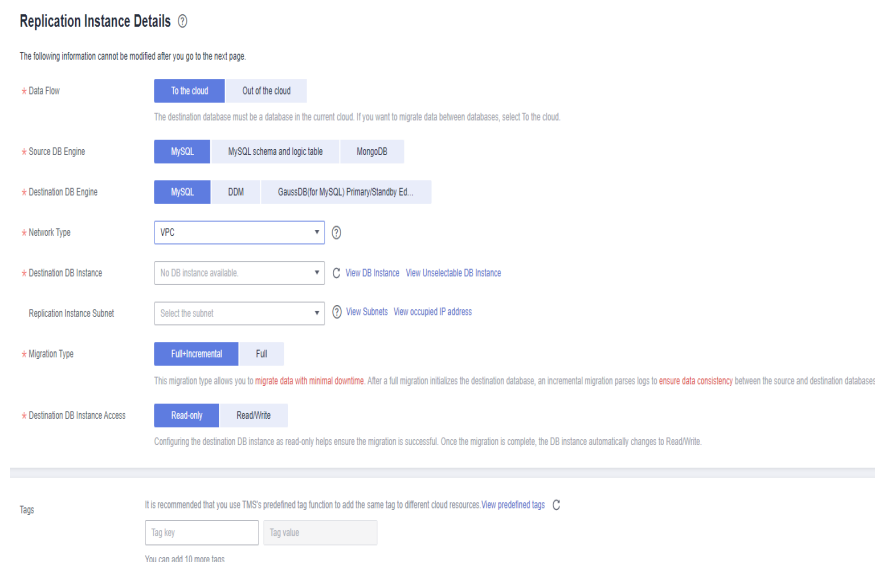
Table 4-2 Replication instance settings

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	Select **VPC**. Available options: **VPC**, **VPN or Direct Connect**, and **Public network**. By default, the value is **Public network**. ● VPC is suitable for migrations of cloud databases. ● Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP. ● VPN is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. ● Direct Connect is suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions. For details about networks, see **Preparations**.
Destination DB Instance	The RDS DB instance you created.
Replication Instance Subnet	The subnet where the replication instance resides. You can also click **View Subnet** to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.
Destination Database Access	● **Read-only** During 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 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. The task cannot be modified after being created.

Parameter	Description
Migration Type	• **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.
Enterprise Project	• If the DB instance has been associated with an enterprise project, select the target project from the **Enterprise Project** drop-down list. • You can also go to the ProjectMan console to create a project. For details about how to create a project, see the *ProjectMan User Guide*.
Tags	• This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. • After a task is created, you can view its tag details on the **Tags** tab. For details, see **Tag Management**.

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**.

📖 **NOTE**

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

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

Figure 4-4 Self-build on ECS - source database information

Table 4-3 Self-build on ECS - source database information

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

Parameter	Description
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. **NOTE** – The maximum size of a single certificate file that can be uploaded is 500 KB. – If the SSL certificate is not used, your data may be at risk.

> 📖 **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 4-5 RDS DB instance - source database information
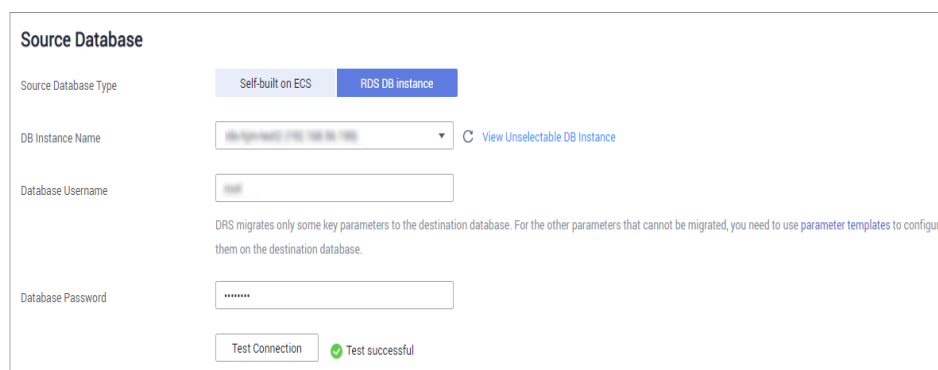
Table 4-4 RDS DB instance - source database information

Parameter	Description
Source 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.

- Destination database configuration

Figure 4-6 Destination database information

Table 4-5 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.
Database Password	The password for the database username.
Migrate Definer to User	– **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?** – **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

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 4-7 Migration type

Table 4-6 Migration types and objects

Parameter	Description

Flow Control	You can choose whether to control the flow.
	
	● **Yes**
	You can customize the maximum migration speed.
	
	In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is **All day**. A maximum of three time ranges can be set, and they cannot overlap.
	
	The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.
	
	Figure 4-8 Flow control
	
	● **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.
	NOTE
	  – 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**.

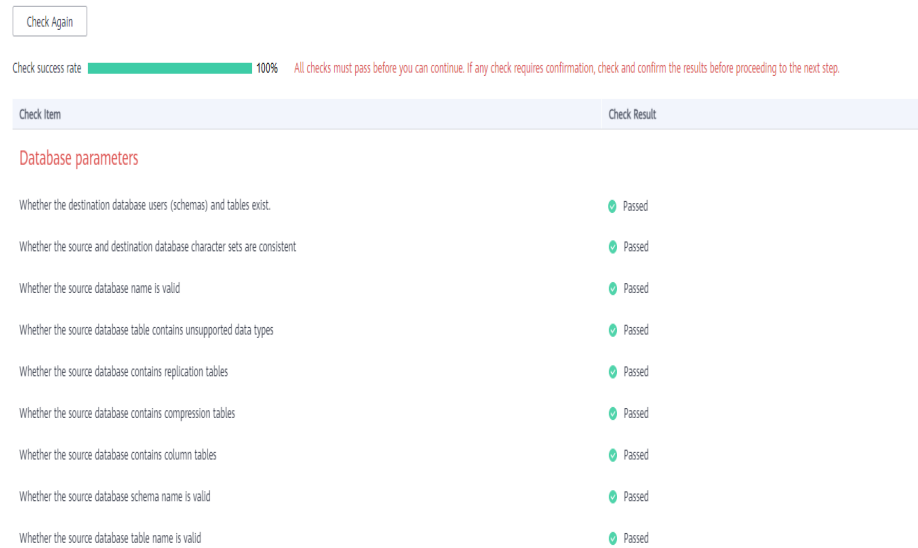| Parameter	Description
Take Snapshot	If you perform a full migration, you can take a snapshot for your databases. ● **No** This option applies to exports for which no data is written to the source database. If data is modified during a full migration, the exported data is point in time inconsistent. The stability and performance of a migration without a snapshot taken is better than that of a migration with a snapshot taken. ● **Yes** A snapshot with consistent data at the point in time is generated during service running. Data changes during migration are not shown in the exported data. **NOTE** – Snapshot reads use MySQL backup lock to lock global tables and automatically unlock them within 3s after consistent reads are enabled. To prevent full migration failures, take a snapshot when the source database is idle and does not perform DML or DDL operations during snapshot migration. – Only MySQL full migration tasks support the snapshot mode. To use this function, you can submit a whitelist application. – Do not perform DDL operations during migration in snapshot mode. Otherwise, full migration will fail.
Migrate Account	During a database migration, accounts need to be migrated separately. 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. ● **Yes** If you need to migrate accounts, see **Migrating Accounts**. ● **No** During migration, accounts, permissions, and passwords are not migrated.
Filter DROP DATABASE	To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations. ● 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
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **NOTE** ● 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. ● 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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 4-9 Pre-check

- 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.
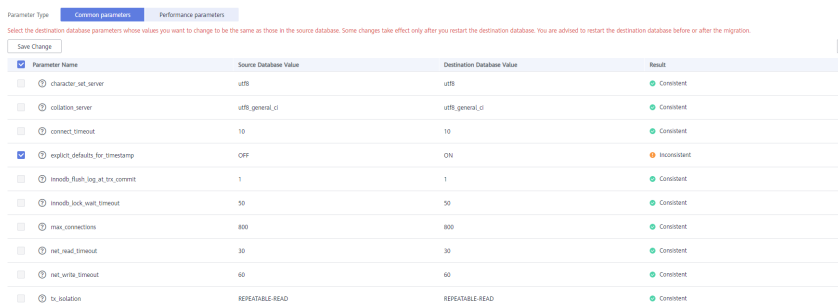
Step 6 Compare source and destination parameters.

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 4-10 Modifying common parameters

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:

i.   Click **Use Source Database Value**.

DRS automatically makes the destination database values the same as those of the source database.

Figure 4-11 One-click modification

📖 **NOTE**

You can also manually enter parameter values.

ii. 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 4-12 Performance parameters

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 Notification**, **SMN Topic**, **Synchronization Delay Threshold**, and **Stop Abnormal Tasks After** and confirm that the configured information is correct and click **Submit** to submit the task.

Figure 4-13 Task startup settings

Start Time	Start upon task creation	Start at a specified time	⑦
Send Notifications	🔵 ⑦ Please handle exceptions within 48 hours of receiving SMS messages or emails.
*SMN Topic	▼ C ⑦
Synchronization Delay Threshold(s)	⚪ ⑦
*Stop Abnormal Tasks After	14	⑦ Abnormal tasks run longer than the period you set (unit: day) will automatically stop.

Table 4-7 Task startup settings

Parameter	Description
Started Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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	SMN topic. This parameter is optional. If an exception occurs during migration, the system will send a notification to the specified recipients.
SMN Topic	This parameter is available only after you enable Send Notifications and create a topic on the SMN console and add a subscriber. For details, see *Simple Message Notification User Guide*.

Parameter	Description
Synchronization Delay Threshold	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
	● 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 Notification**.
	● If the delay threshold is set to 0, no notifications will be sent to the recipient.
Stop Abnormal Tasks After	Number of days after which an abnormal task is automatically stopped. The value must range from 14 to 100. The default value is **14**.
	NOTE
	Tasks in the abnormal state are still charged. If tasks remain in the abnormal state for a long time, they cannot be resumed. Abnormal tasks run longer than the period you set (unit: day) will automatically stop to avoid unnecessary fees.

Step 8 After the task is submitted, 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

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.

● With the progress bar, you can view the migration progress of structures, data, and indexes. When the progress reaches 100%, the migration is complete. The migration of data and indexes is relatively slow.

● 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 only view details about MySQL to MySQL, MongoDB to DDS, and DDS to MongoDB migration 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 4-14 Migration progress overview
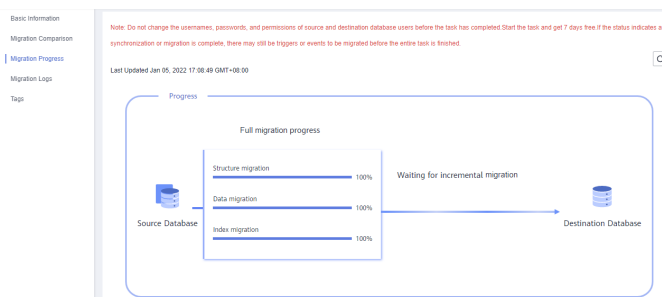
  You can also view the incremental migration delay on the **Online Migration Management** page. When the incremental migration delay exceeds the preset or default threshold, the value of the incremental migration delay is displayed in red in the task list.

   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. After the incremental migration starts, the progress is not displayed. You can click the **Migration Comparison** tab to compare the data consistency.

Figure 4-15 Migration progress details

- 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 4-16 Performance monitoring
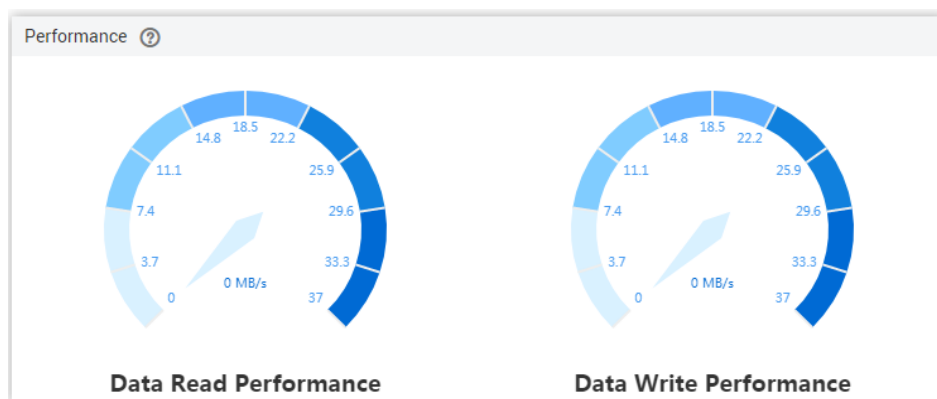
----End

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 4-17 Viewing migration logs

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

----End

4.4 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 4-18 Migration comparison process
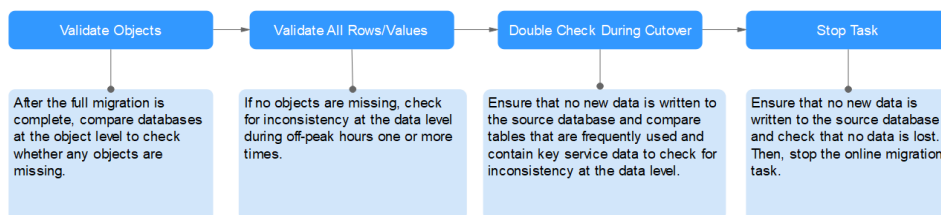
DRS supports object-level, account-level, and data-level comparisons.

- Object-level comparison: compares objects such as databases, indexes, tables, views, stored procedures, functions, and table sorting rules.
- Data-level comparison: compares rows and values of tables or collections.

📖 **NOTE**

- Full migration tasks do not support data-level comparisons.
- If DDL operations were performed on the source database, you need to compare the objects again to ensure the accuracy of the comparison results.
- Some data types do not support value comparison. For details, see **Which Data Types Does Not Support Content Comparison?**
- To prevent resources from being occupied for a long time, DRS limits the row comparison duration. If the row comparison duration exceeds the threshold, the row comparison task stops automatically. If the source database is a relational database, the row comparison duration is 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration is 30 minutes.

- Account comparison: compares the account names and permissions of the source and destination databases.

Table 4-8 Supported comparison mode

Migration Direction	Data Flow	Object-level Comparison	Row Comparison	Value Comparison	Account-level Comparison
To the cloud	MySQL->MySQL	Yes	Yes	Yes	Yes
To the cloud	MySQL -> GaussDB(for MySQL) primary/standby	Yes	Yes	Yes	No
To the cloud	MySQL->DDM	Yes	Yes	No	No
To the cloud	MongoDB->DDS	Yes	Yes	Yes	Yes
To the cloud	MongoDB->GaussDB(for Mongo)	Yes	Yes	Yes	No
To the cloud	MySQL schema and logic table -> DDM	Yes	Yes	No	No
From the cloud	MySQL->MySQL	Yes	Yes	Yes	Yes
From the cloud	DDS->MongoDB	Yes	Yes	Yes	No

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.
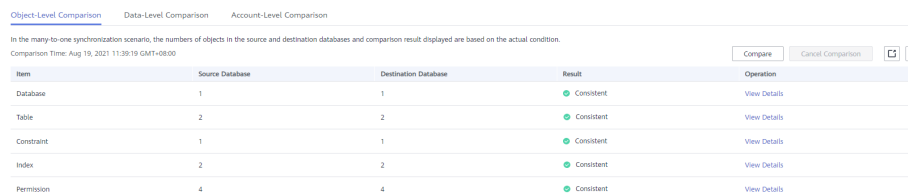
1. Check the integrity of the database object.

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

   Figure 4-19 Comparing objects

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

2. After the check is complete, compare the number of rows and values.

   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.

   a. In the **Before You Start** pane, click **Validate All Rows/Values**.

   b. In the displayed **Create Comparison Task** dialog box, specify **Comparison Type**, **Comparison Method**, **Comparison Time**, and **Object**. Then, click **OK**.

Figure 4-20 Creating a comparison task
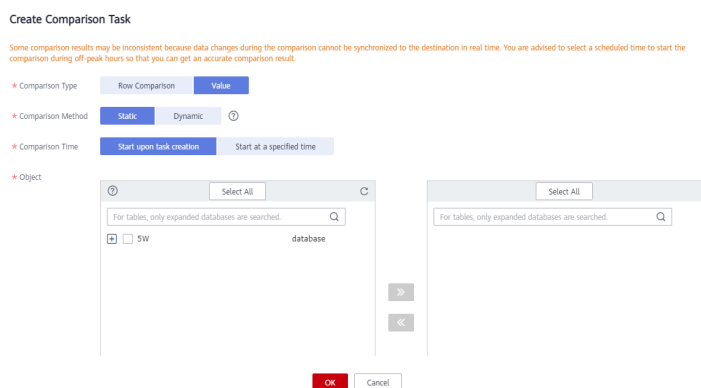
- **Comparison Type**: compares rows and values.

- **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.

- **Comparison Time**: You can select **Start upon task creation** or **Start at a specified time**. There is a slight difference in time between the source and destination databases during synchronization. Data inconsistency may occur. You are advised to compare migration items during off-peak hours for more accurate results.

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

  ☐ **NOTE**

  - Before the value comparison, the system will help you evaluate the time required for value comparison.

  - When you select an object, the spaces before and after the object name are not displayed. If there are multiple spaces in the middle of the object name, only one space is displayed.

After the comparison creation task is submitted, the **Data-Level Comparison** tab is displayed. Click ⟳ to refresh the list and view the comparison result of the specified comparison type.

Figure 4-21 Viewing the data-level comparison result

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.

Figure 4-22 Viewing comparison details

> **☐☐ NOTE**
>
> You can cancel a running task at any time and view the comparison report of a canceled comparison task.

3. Compare database accounts and permissions. Click the **Account-Level Comparison** tab to view the comparison results of database accounts and permissions.

Figure 4-23 Account-level comparison

> **☐☐ NOTE**
>
> – Full migration tasks do not support account-level comparisons.

4. Perform a double check before the cutover.

In the **Before You Start** pane, 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 2.2**.

5. 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

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, accounts to be migrated can be classified into the following types: 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 4-24 Changing the IP address
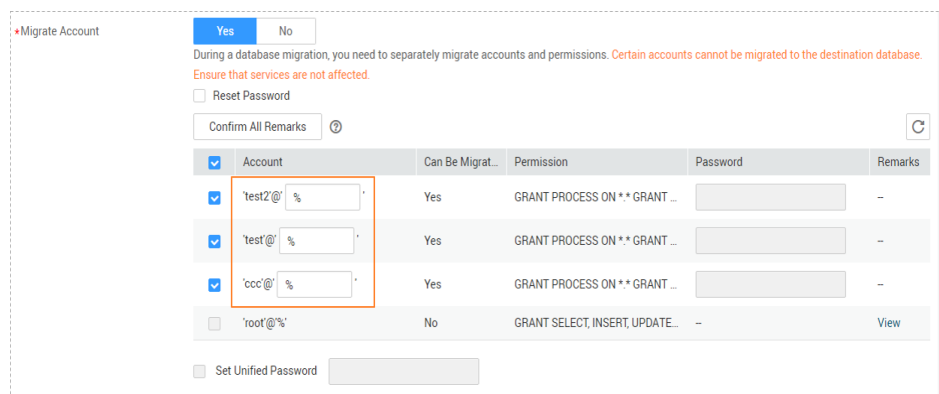
Step 2 By default, account permissions cannot be modified. For accounts that can be migrated (including accounts that can be completely migrated and accounts whose permissions need to be reduced), the system also migrates the permissions of these accounts.

After the migration is successful, accounts in the destination database are those whose permissions need to be reduced.
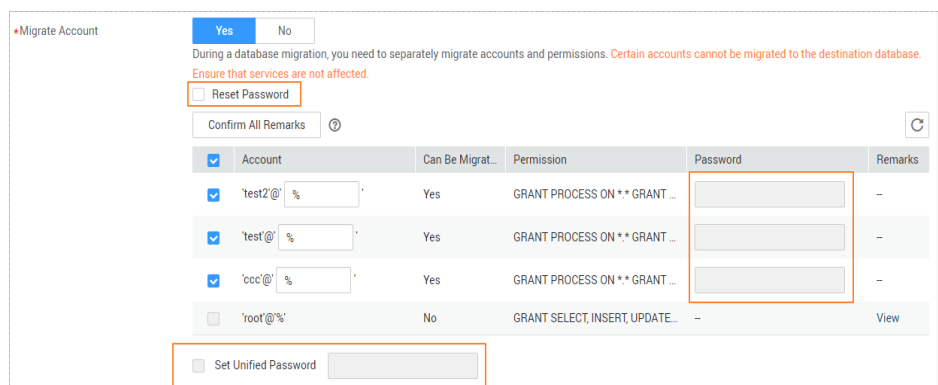
Figure 4-25 Account permissions
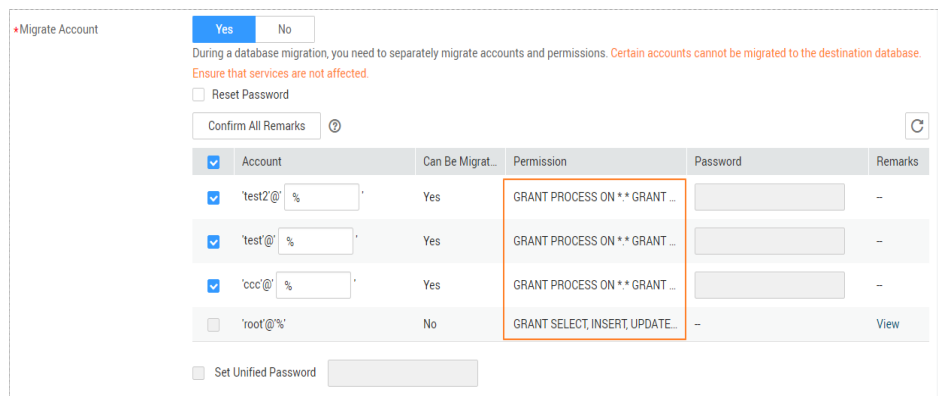
Step 3 Migrate account passwords.

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 4-26 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 4-27 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 4-28 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

- Currently, only MySQL supports account migration.
- 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 4-29 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 4-30 Remarks
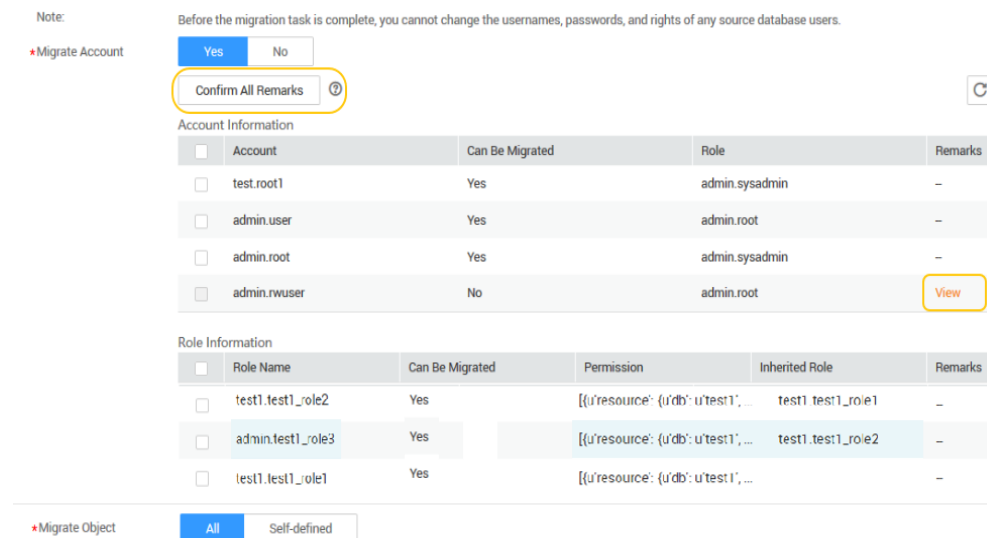
----End

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 4-9 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

Parameter	Type	Restart Required
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_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	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
sync_binlog	Performance parameter	No

MySQL 5.7

Table 4-10 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

Parameter	Type	Restart Required
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_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
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

MySQL 8.0

Table 4-11 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

Parameter	Type	Restart Required
net_write_timeout	Common parameter	No
explicit_defaults_for_time stamp	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
collation_server	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
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.

4.6 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 Editing Migration 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 4-12 Task information

Task Information	Description
Task Name	The task name consists of 4 to 50 characters, starts with a letter, and can contain only letters (case-insensitive), 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.
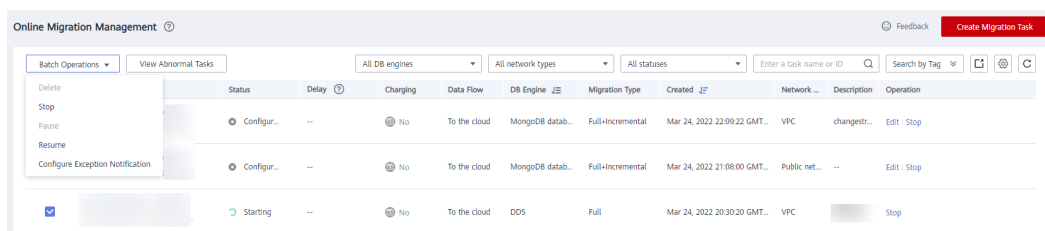
----End

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 4-31 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 Migration Information

During the migration, you may change the password of the source or destination database. As a result, the migration task fails. In this case, you need to change the password on the DRS console and resume the task.

You can modify the following information:

- Source database password

- Destination database password

☐ NOTE

> After the preceding information is changed, the change takes effect immediately, and the data in the destination database is not cleared.

Prerequisites

You have logged in to the DRS console.

Procedure

Step 1 On the **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 passwords of the source and destination databases and click **OK**.

Step 4 View the change result on the **Basic Information** tab.

----End

4.6.4 Editing a Migration Task

This section describes how to modify configuration information of a migration task, including information about the task, replication instance, and migration. For migration tasks in the following statuses, you can edit the tasks again after the replication instances are created:

- Creating

- Configuration

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 4-32 Migration type

Table 4-13 Migration types and objects

Parameter	Description

Flow Control	You can choose whether to control the flow.
	
	● **Yes**
	You can customize the maximum migration speed.
	
	In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is **All day**. A maximum of three time ranges can be set, and they cannot overlap.
	
	The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.
	
	Figure 4-33 Flow control
	
	● **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.
	
	NOTE
	
	   – 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

Take Snapshot	If you perform a full migration, you can take a snapshot for your databases.
	
	● **No**
	This option applies to exports for which no data is written to the source database. If data is modified during a full migration, the exported data is point in time inconsistent. The stability and performance of a migration without a snapshot taken is better than that of a migration with a snapshot taken.
	
	● **Yes**
	A snapshot with consistent data at the point in time is generated during service running. Data changes during migration are not shown in the exported data.
	
	NOTE
	  – Snapshot reads use MySQL backup lock to lock global tables and automatically unlock them within 3s after consistent reads are enabled. To prevent full migration failures, take a snapshot when the source database is idle and does not perform DML or DDL operations during snapshot migration.
	  – Only MySQL full migration tasks support the snapshot mode. To use this function, you can submit a whitelist application.
	  – Do not perform DDL operations during migration in snapshot mode. Otherwise, full migration will fail.

Migrate Account	During a database migration, accounts need to be migrated separately. 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. ● **Yes** If you need to migrate accounts, see **Migrating Accounts**. ● **No** During migration, accounts, permissions, and passwords are not migrated.
Filter DROP DATABASE	To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations. ● 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
Migrate Object	You can choose to migrate all objects, tables, or databases based on your service requirements. ● **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. 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. **NOTE** ● 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. ● 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. ● The search function can help you quickly select the required database objects.

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 **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.

Figure 4-34 Pre-check
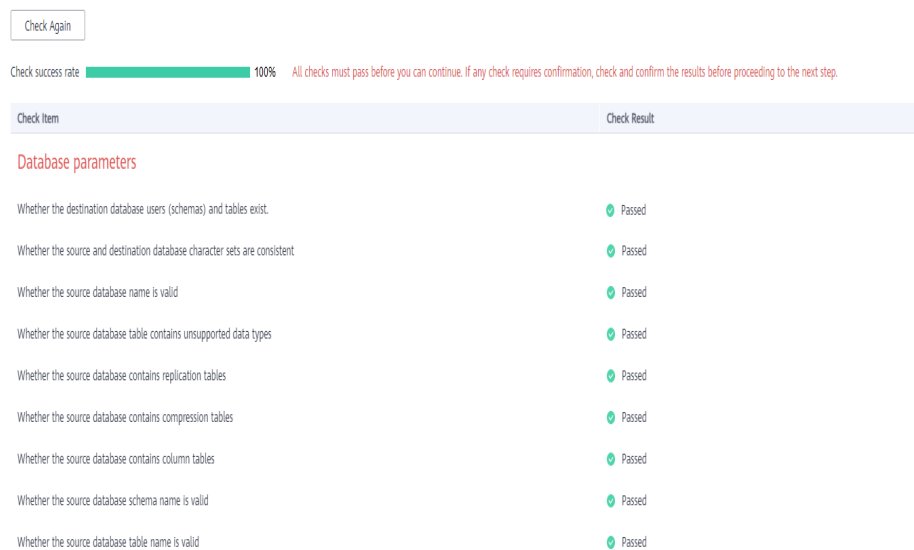
- If the check is complete and the check success rate is 100%, click **Next**.

  ◫ **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

4.6.5 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

📖 **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.
- If the full migration is performed in snapshot mode and fails, it cannot be resumed.

Prerequisites

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

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 left corner.
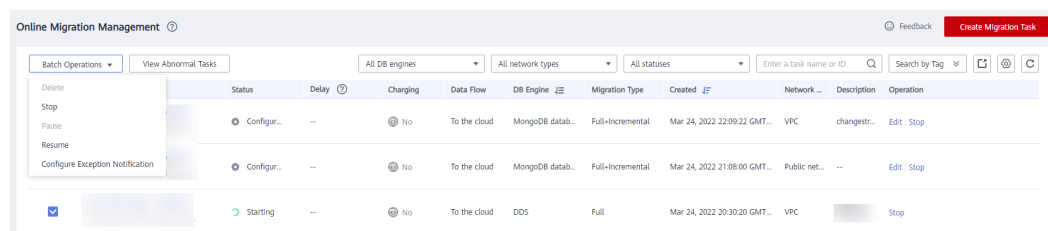
----End

Resume Tasks

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

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

Figure 4-35 Batch Operations

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

----End

4.6.6 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 failed migration tasks in any of the following statuses:

- Migration failure status

You can reset the following DRS tasks:

- MySQL->MySQL
- MySQL->DDM
- MySQL -> GaussDB(for MySQL) primary/standby
- MySQL schema and logic table -> DDM
- MongoDB->DDS
- DDS->MongoDB

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 left corner.

Step 3 Perform **Step 2** to **Step 3** from method 1.

----End

4.6.7 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) primary/standby
  - MongoDB->DDS
- 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

- After the task is paused, the status of the task becomes **Paused**.
- After you select **Pause log capturing**, the DRS instance will no longer communicate with the source and destination databases. If the pause duration is too long, the task may fail to be resumed because the logs required by the source database expire. It is recommended that the pause duration be less than or equal to 24 hours.
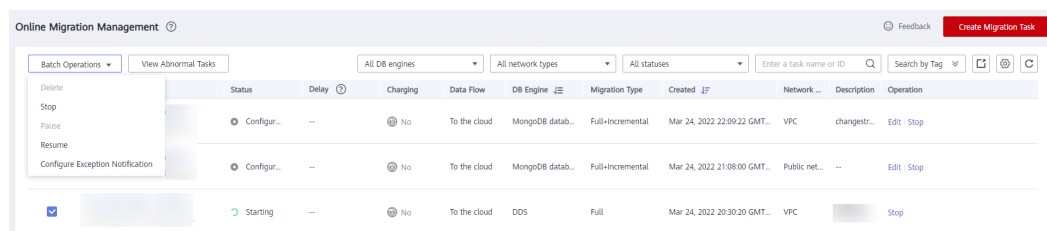- You can use the resumable transfer function to continue the migration.

----End

Pausing Tasks

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

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

Figure 4-36 Batch Operations

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

----End

4.6.8 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 links support this function:

- To the cloud
  - MySQL->MySQL
  - MySQL -> GaussDB(for MySQL) primary/standby
- From of the cloud
  - MySQL->MySQL

📖 NOTE

- Flow control mode takes effect only during a full migration.
- 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 **Migration Information** area on the **Basic Information** tab, click **Modify** next to the **Flow Control** field.

Figure 4-37 Basic information
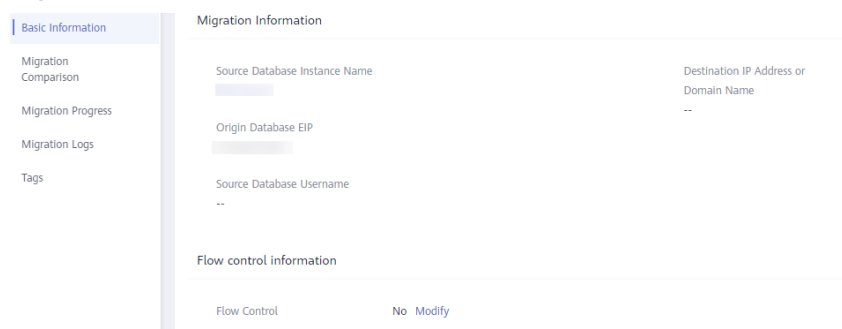
Step 2 In the displayed dialog box, modify the settings.

Figure 4-38 Modify Flow Control
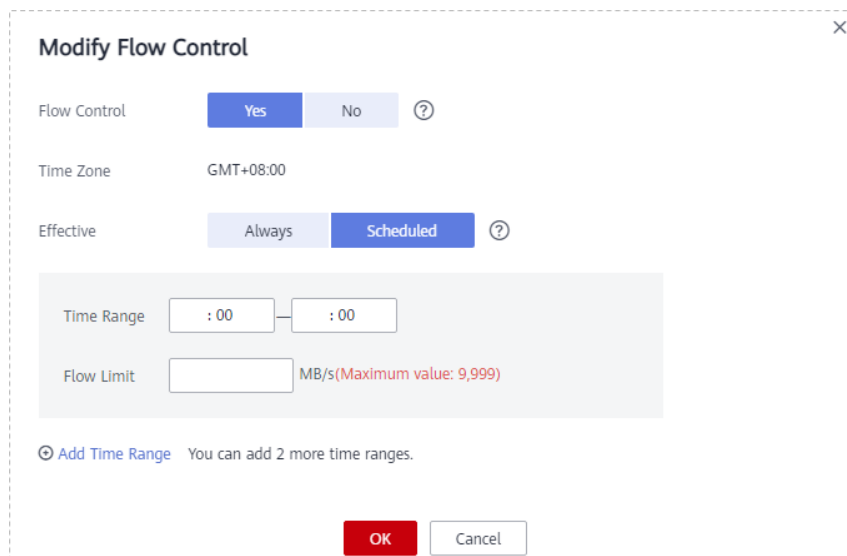
    ----End

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 4-39 Task list

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

    Figure 4-40 Modify Flow Control
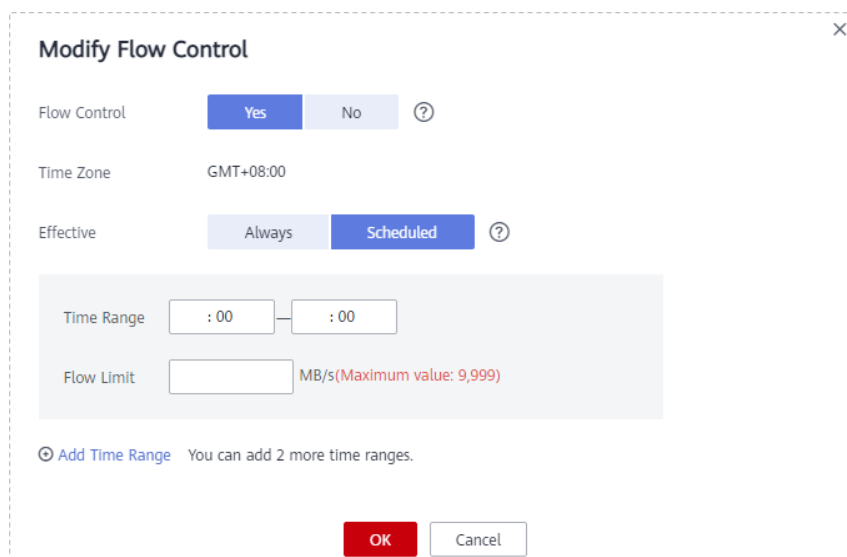
    ----End

4.6.9 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.

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 **OK**.

📖 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.

----End

Stopping Tasks

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

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

Figure 4-41 Batch Operations
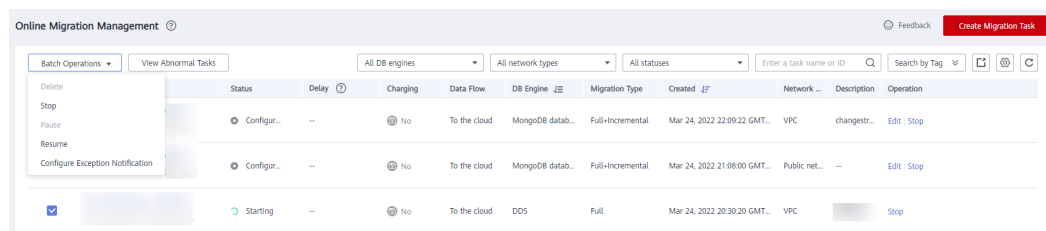
Step 3　In the displayed dialog box, confirm the task information and click **Yes**.

----End

4.6.10 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 4-42 Batch Operations
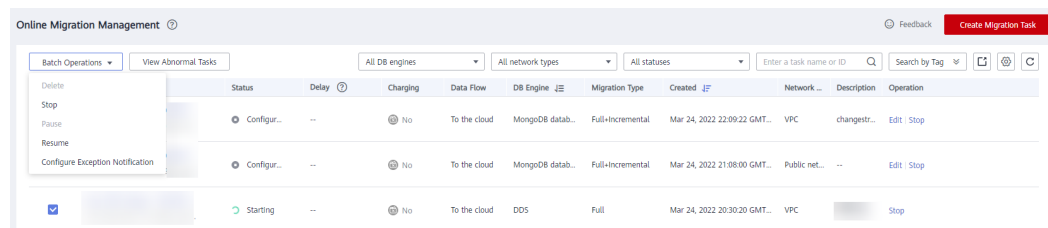
> **Step 3** In the displayed dialog box, confirm the task information and click **Yes**.

> **----End**

4.6.11 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 colne 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

4.6.12 Task Statuses

Migration statuses indicate different migration phases.

Table 4-14 lists statuses and descriptions of online migration tasks.

Table 4-14 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.
Configuring	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 starting.
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.
Fault rectification	A replication instance is faulty and the system automatically restores the migration task.
Cloning	A migration task is being cloned.

Status	Description
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

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 10 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 On the **Basic Information** tab, click the **Tags** tab.

Step 3 On the **Tags** tab, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.

- 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 and can contain 1 to 128 characters, including letters, digits, spaces, and special characters _:=+.-@

- The tag value can be empty. It cannot start or end with a space and can contain 0 to 255 characters, including letters, digits, spaces, and special characters _:=+.-@

Step 4 After a tag has been added, you can view and manage it on the **Tags** page.

----End

Editing a Tag

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

Step 2 On the **Basic Information** tab, click the **Tags** tab.

Step 3 On the **Tags** page, click **Add/Edit Tags**. In the displayed dialog box, modify the tag and click **OK**.

----End

Delete a Tag

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

Step 2 On the **Basic Information** tab, click the **Tags** tab.

Step 3 On the **Tags** page, locate the tag to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

Step 4 After the tag is deleted, it will no longer be displayed on the **Tags** page.

----End

6 Interconnecting with CTS

6.1 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 6-1 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

6.2 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

7 Interconnecting with Cloud Eye

7.1 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 7-1 lists the DRS performance metrics.

Table 7-1 DRS metrics

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
cpu_util	CPU Usage	CPU usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
mem_util	Memory Usage	Memory usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
network_incoming_bytes_rate	Network Input Throughput	Incoming traffic in bytes per second	≥ 0 bytes /s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
network_outgoing_bytes_rate	Network Output Throughput	Outgoing traffic in bytes per second	≥ 0 bytes /s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
disk_read_bytes_rate	Disk Read Throughput	Number of bytes read from the disk per second (bytes/second).	≥ 0 bytes /s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
disk_write_bytes_rate	Disk Write Throughput	Number of bytes written to the disk per second (bytes/second).	≥ 0 bytes/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
disk_util	Storage Space Usage	Storage space usage of the monitored object	0-100%	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_bytes_rate	Source Database Read Throughput	Table data or WAL bytes read from the source database per second	≥ 0 bytes/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_rows_rate	Rows Read from Source Database per Second	Number of table data rows or WAL rows read from the source database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
extract_latency	Source Database WAL Extract Lag	Latency of extracting WAL from the source database Unit: ms.	≥ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_bytes_rate	Destination Database Write Throughput	Number of bytes written to the destination database per second.	≥ 0 bytes/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_rows_rate	Rows Written into Destination Database per Second	Number of rows that are written to the destination database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_transactions_rate	DML TPS	Number of DML transactions written to the destination database per second.	≥ 0 transaction/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_ddls_rate	DDL TPS	Number of DDLs written to the destination database per second.	≥ 0 transaction/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_latency	Replication Delay	Delay (in milliseconds) of data replay.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_average_execute_time	Average Transaction Execution Time	Average execution time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is millisecond.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_average_commit_time	Average Transaction Commit Time	Average commit time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is ms.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_current_state	Synchronization Status	This metric is the synchronization status of the current kernel data (10: abnormal; 1: idle; 2: DML; 3: DDL), instead of the task status.	10: abnormal 1: idle 2: DML is executed. 3: DDL is executed.	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_thread_workers	Synchronization Threads	Number of working threads for data synchronization	≥ 0	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

Dimensions

Key	Value
instance_id	DRS instance ID

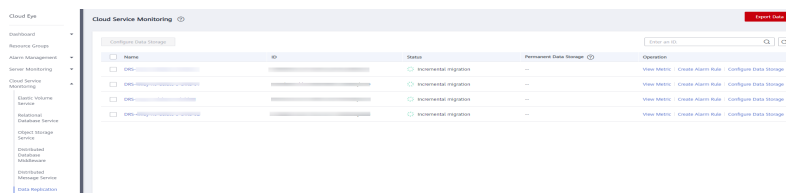# 7.2 Configuring Alarm Rules

Scenarios

You can configure DRS alarm rules to customize the monitored objects and notification policies and learn the DRS running status in a timely manner.

This section describes how to set DRS alarm rules, including the alarm rule name, service, dimension, monitoring scope, template, and whether to send a notification.

Procedure

Step 1 Log in to the management console.

Step 2 Under **Management & Governance**, click **Cloud Eye**.

Step 3 In the navigation pane on the left, choose **Cloud Eye** > **Data Replication Service**.

Figure 7-1 Choosing a monitored object

Step 4 Select the DB instance which you want to create an alarm rule for and click **Create Alarm Rule** in the **Operation** column.

Step 5 On the displayed page, set parameters as required.

Figure 7-2 Configuring alarm information

- 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

7.3 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 30 days.

Figure 7-3 Viewing monitoring metrics
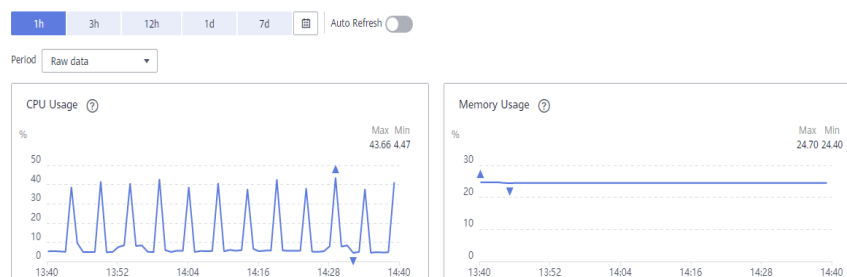
----End

A Change History

Released On	Description
2022-09-30	This issue is the first official release.