Data Replication Service

Best Practices

Issue 01

Date 2022-09-30

HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2022. All rights reserved.

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

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Huawei Technologies Co., Ltd.

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

Website: https://www.huawei.com

Email: support@huawei.com

Contents

 Copyright © Huawei Technologies Co., Ltd.

 Copyright © Huawei Technologies Co., Ltd.

1 Best Practices Summary

This document describes how to use Data Replication Service (DRS) to quickly migrate or synchronize databases in typical application scenarios.

Table 1-1 DRS best practices

Function	Source DB Type	Destination DB	Cases
Real-Time Migration	MySQL databases on other clouds	Huawei Cloud GaussDB(for MySQL) instances	From Other Cloud MySQL to GaussDB(for MySQL)
	MongoDB databases on other clouds	Huawei Cloud DDS	From Other Cloud MongoDB to DDS
	MongoDB databases on ECSs	Huawei Cloud DDS	From MongoDB on ECS to DDS
	On-premises MySQL databases	Huawei Cloud RDS MySQL instances	From On-Premises MySQL to RDS MySQL
	On-premises MongoDB databases	Huawei Cloud DDS	From On-Premises MongoDB to DDS
Backup Migration	On-premises Microsoft SQL Server databases	Huawei Cloud RDS SQL Server instances	Migrating Microsoft SQL Server Backup Data to RDS SQL Server DB Instance
Real-Time Synchronization	PostgreSQL databases on other clouds	Huawei Cloud RDS PostgreSQL instances	From Other Cloud PostgreSQL to RDS PostgreSQL

 Copyright © Huawei Technologies Co., Ltd.

Function	Source DB Type	Destination DB	Cases
	PostgreSQL databases on ECSs		From PostgreSQL on ECS to RDS PostgreSQL
	On-premises PostgreSQL databases		From On-Premises PostgreSQL to RDS PostgreSQL
	RDS for MySQL instances	DMS for Kafka	From RDS MySQL to Kafka

 Copyright © Huawei Technologies Co., Ltd.

2 Real-Time Migration

2.1 From Other Cloud MySQL to GaussDB(for MySQL)

2.1.1 Overview

Description

This section includes the following content:

- Create a GaussDB(for MySQL) instance.
- Migrate data from MySQL on other clouds to GaussDB(for MySQL).

Prerequisites

- You have registered with Huawei Cloud.
- Your account balance is greater than or equal to $0 USD.

Deployment Architecture

In this example, the source is a MySQL database on other cloud platforms and the destination is a Huawei Cloud GaussDB(for MySQL) instance. Data is migrated from the source to the destination over a public network. For details about the deployment architecture, see **Figure 2-1**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-1 Deployment architecture

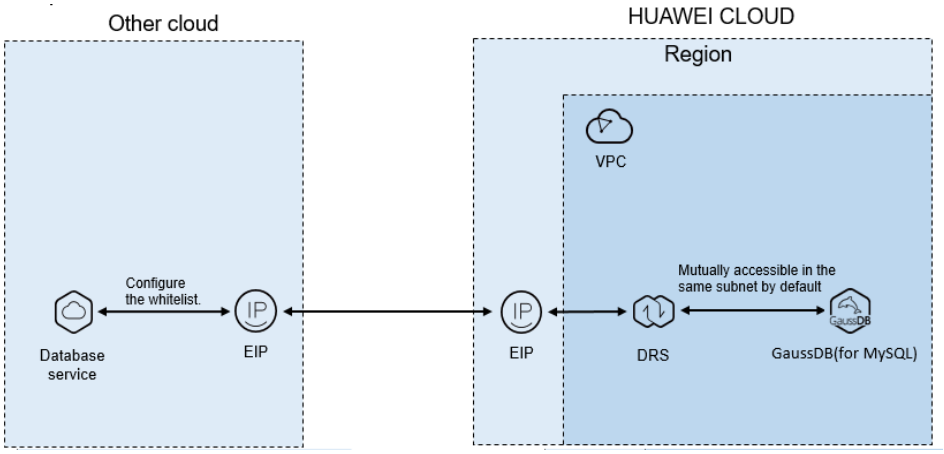

Service List

- Virtual Private Cloud (VPC)
- GaussDB(for MySQL)
- Data Replication Service (DRS)

Before You Start

- The resource planning in this best practice is for demonstration only. Adjust it as needed.
- All settings in this best practice are for reference only. For more information about MySQL migration, see **From MySQL to GaussDB(for MySQL) Primary/Standby**.

2.1.2 Resource Planning

Table 2-1 Resource planning

Category	Subcategory	Plan	Description
VPC	VPC name	vpc-DRStest	Specify a name that is easy to identify.
	Region	EU-Dublin	To achieve lower network latency, select the region nearest to you.
	AZ	AZ 1	-
	Subnet	10.0.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-drs01	Specify a name that is easy to identify.

 Copyright © Huawei Technologies Co., Ltd.

Category	Subcategory	Plan	Description
Other cloud MySQL	DB engine version	MySQL 5.7	-
	IP address	10.154.217.42	Enter an IP address.
	Port	3306	-
GaussDB(for MySQL) instance	Instance name	gauss-drstar	Specify a name that is easy to identify.
	DB engine version	MySQL 8.0	-
	AZ type	Single AZ	In this example, a single AZ is used. To improve service reliability, select multiple AZs.
	AZ	AZ1	AZ1 is selected in this example. To improve service reliability, deploy the instance across multiple AZs.
	Instance class	Dedicated 4 vCPUs | 16 GB	-
DRS migration task	Task name	DRS-test-migrate	Specify a name that is easy to identify.
	Source DB engine	MySQL	-
	Destination DB engine	GaussDB(for MySQL) primary/ standby instance	-
	Network type	Public network	Public network is used in this example.

 Copyright © Huawei Technologies Co., Ltd.

2.1.3 Operation Process

Figure 2-2 Flowchart

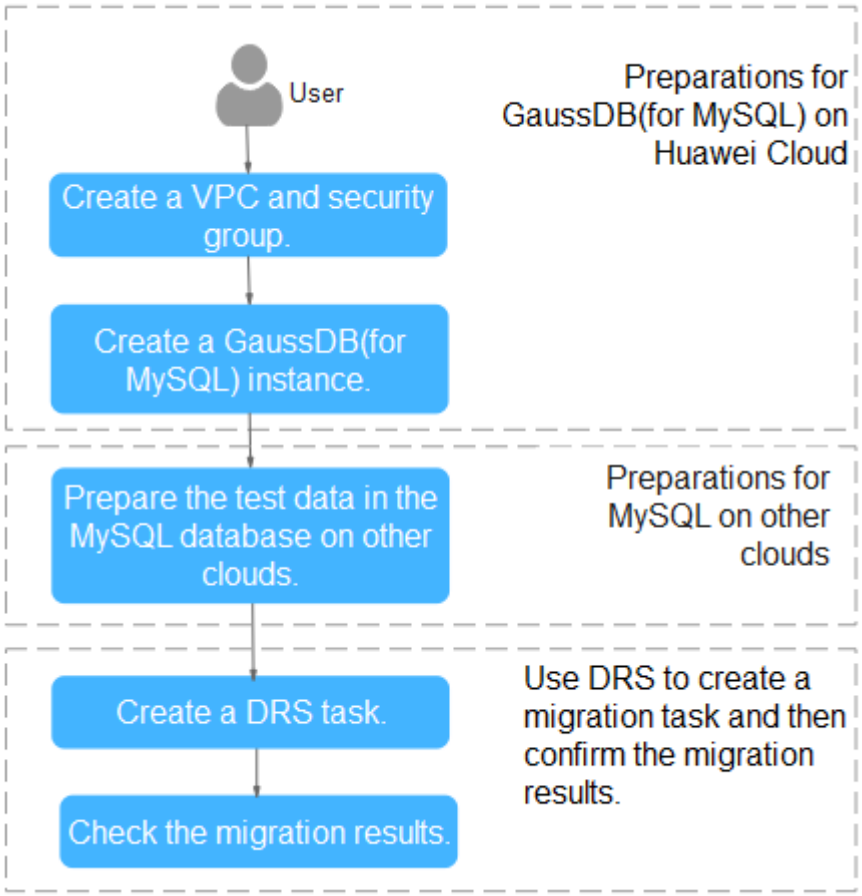

2.1.4 Creating a VPC and Security Group

Create a VPC and security group for a GaussDB(for MySQL) instance.

Creating a VPC

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select region EU-Dublin.

Step 3 Under the service list, choose **Networking** > **Virtual Private Cloud**.

 Copyright © Huawei Technologies Co., Ltd.

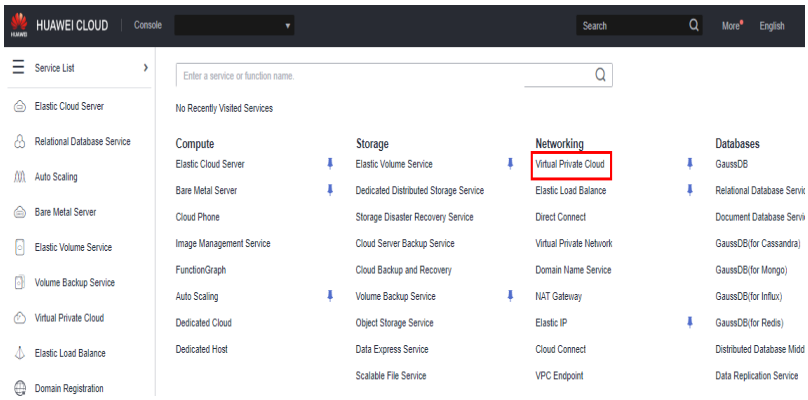

Step 4 Click **Create VPC**.

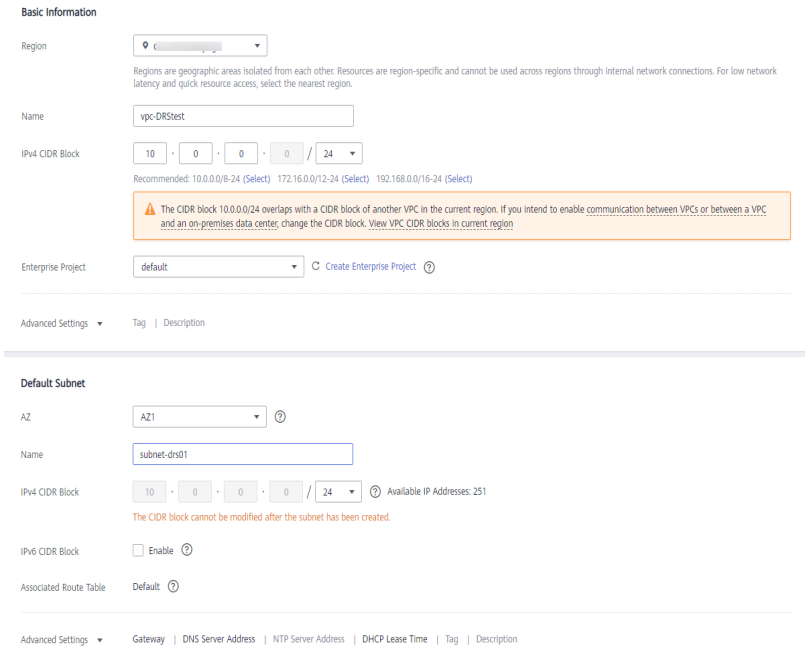

Step 5 Configure parameters as needed and click **Create Now**.

Step 6 Return to the VPC list and check whether the VPC is created.

If the VPC status becomes available, the VPC has been created.

----End

Creating a Security Group

Step 1 Log in to the **management console**.

Step 2 Click  in the upper left corner of the management console and select region EU-Dublin.

Step 3 Under the service list, choose **Networking** > **Virtual Private Cloud**.

 Copyright © Huawei Technologies Co., Ltd.

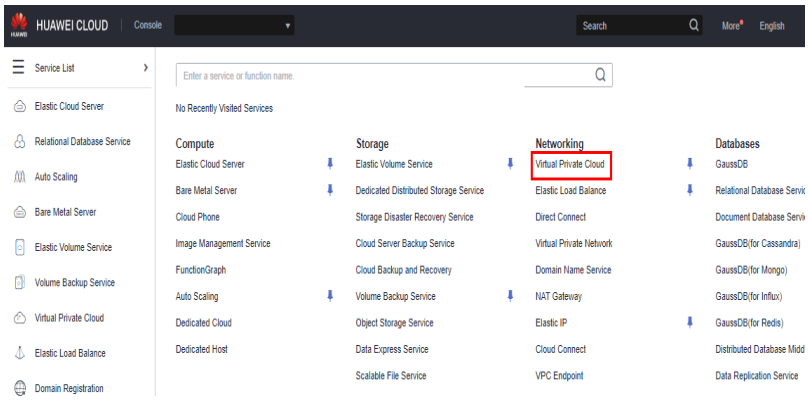

Step 4 In the navigation pane, choose **Access Control** > **Security Groups**.

Step 5 Click **Create Security Group**.

Step 6 Configure parameters as needed.

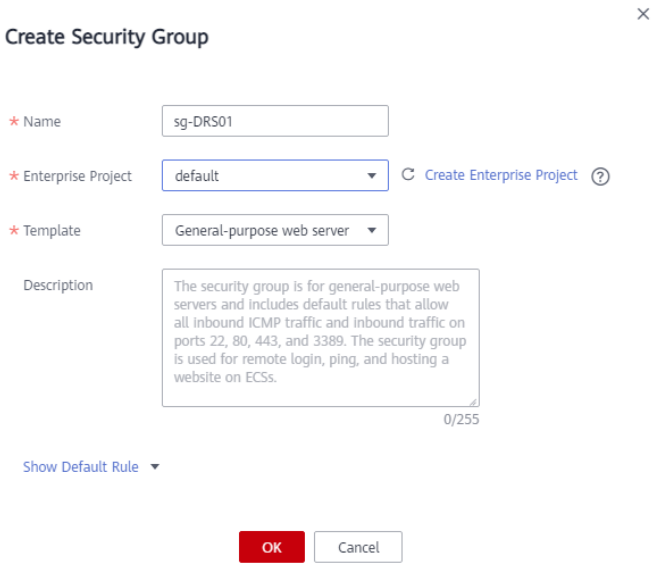

Step 7 Click **OK**.

Step 8 Return to the security group list and click the security group name (**sg-DRS01** in this example).

Step 9 Click the **Inbound Rules** tab, and then click **Add Rule**.

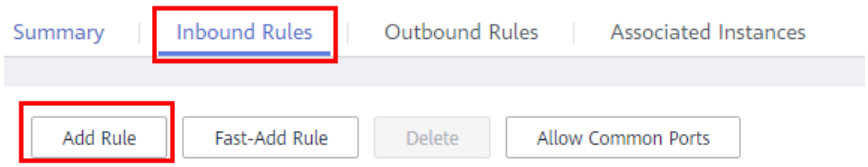

Step 10 Configure an inbound rule to allow access from database port **3306**.

 Copyright © Huawei Technologies Co., Ltd.

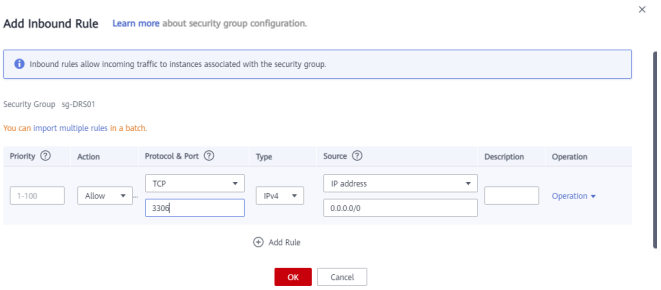

----End

2.1.5 Creating a GaussDB(for MySQL) Instance

This section describes how to create a Huawei Cloud GaussDB(for MySQL) instance.

Step 1 Log in to the **management console**.

Step 2 Click  in the upper left corner of the management console and select region EU-Dublin.

Step 3 Under the service list, choose **Databases** > **GaussDB**.

Step 4 Choose **GaussDB(for MySQL)** and click **Buy DB Instance**.

Step 5 Configure the instance name and basic information.

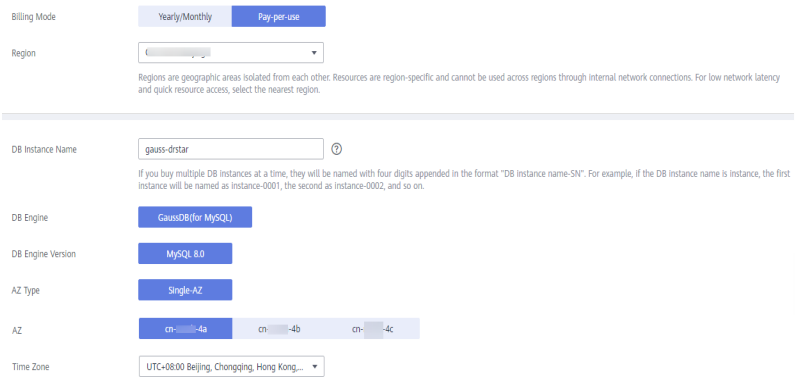

Step 6 Configure instance specifications.

 Copyright © Huawei Technologies Co., Ltd.

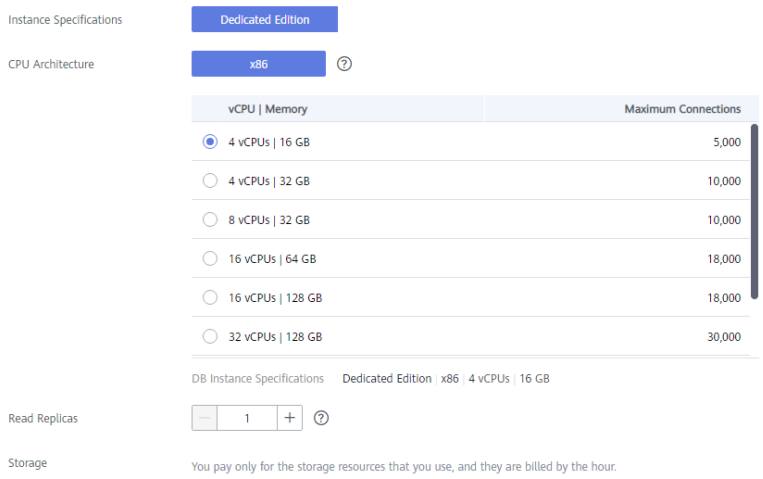

Step 7 Select a VPC and security group for the instance and configure the database port.

The VPC and security group have been created in **Creating a VPC and Security Group**.

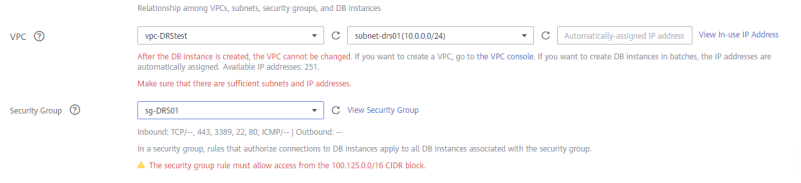

Step 8 Configure the instance password.

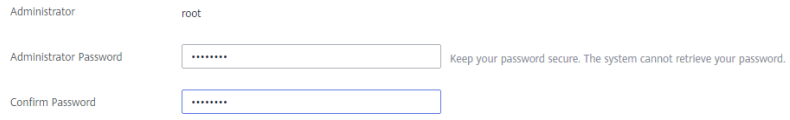

Step 9 Click **Next**. If you do not need to modify your settings, click **Submit**.

Step 10 Return to the instance list. If the instance becomes **Available**, the instance has been created.

----End

2.1.6 Configuring a MySQL Instance on Other Clouds

Prerequisites

- You have purchased a MySQL instance on other platforms.
- The MySQL account has the migration permissions listed in **Permission Requirements**.

Permission Requirements

To migrate data from a MySQL database on other clouds to a GaussDB(for MySQL) instance, the following permissions are required.

 Copyright © Huawei Technologies Co., Ltd.

Table 2-2 Required permissions

Database	Full Migration Permission	Full+Incremental Migration Permission
Source DB (MySQL)	SELECT, SHOW VIEW, and EVENT	SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT

For details about MySQL authorization operations, see **operation guide**.

Network Configuration

Enable public accessibility for the source database. The method for enabling public accessibility depends on the cloud database vendor. For details, see the official documents of the corresponding cloud database vendor.

Take ApsaraDB RDS for MySQL as an example. You need to apply for an external IP address to allow interconnection with external applications. For details about the operations and precautions, see the official documents.

2.1.7 Creating a DRS Migration Task

This section describes how to create a DRS instance and migrate data from a MySQL database on other clouds to a GaussDB(for MySQL) instance.

Pre-migration Check

Before creating a migration task, check the migration conditions.

This section describes how to migrate data from a MySQL database to GaussDB(for MySQL). For details, see **Before You Start**.

Creating a Migration Task

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select region EU-Dublin.

Step 3 Under the service list, choose **Databases** > **Data Replication Service**.

Step 4 In the upper right corner, click **Create Migration Task**.

Step 5 Configure parameters as needed.

1. Specify a migration task name.

 Copyright © Huawei Technologies Co., Ltd.

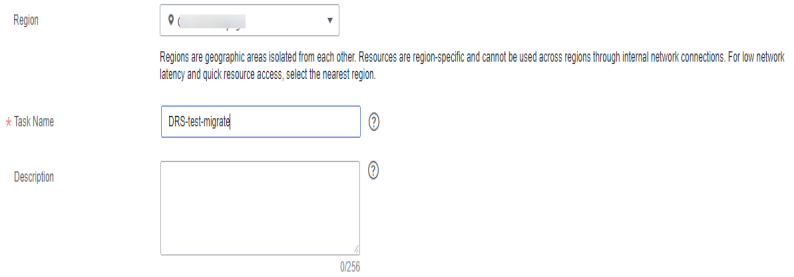

2. Configure replication instance details as needed.

   Select the GaussDB(for MySQL) instance created in **Creating a GaussDB(for MySQL) Instance** as the destination database.

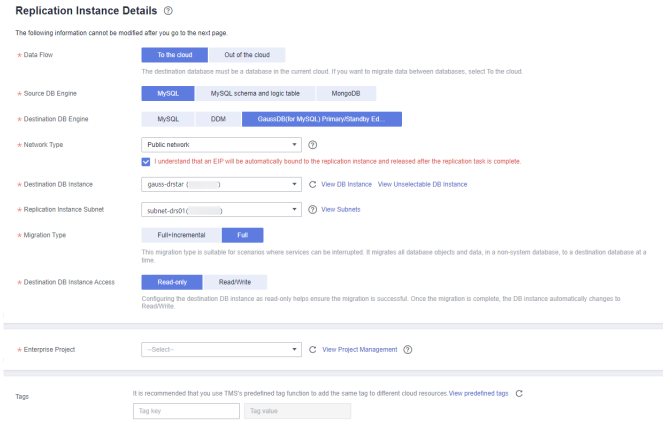

Step 6 Click **Next**.

It takes about 5 to 10 minutes to create a replication instance.

Step 7 Configure a whitelist for the source database to manage network access.

Add the EIP of the DRS replication instance to the whitelist of the source MySQL database to ensure that the source database can communicate with the DRS instance.

The method for configuring the whitelist depends on the cloud database vendor. For details, see the official documents of the corresponding cloud database vendor.

Step 8 Configure source and destination database information.

1. Configure the source database information and click **Test Connection**. If a successful test message is returned, login to the destination is successful.

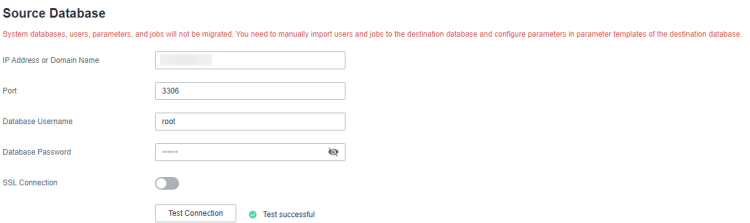

2. Configure destination database information and click **Test Connection**. If a successful test message is returned, login to the destination is successful.

 Copyright © Huawei Technologies Co., Ltd.

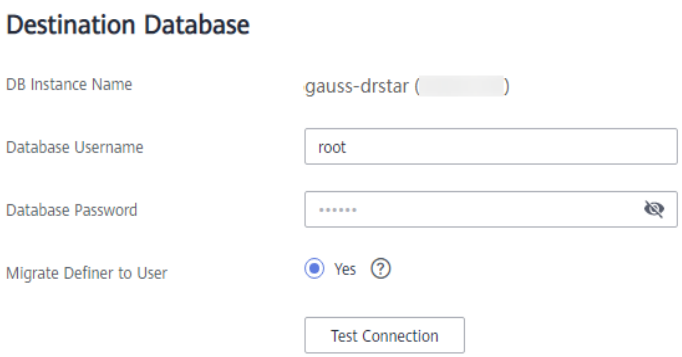

Step 9 Click **Next**.

Step 10 On the **Set Task** page, select migration accounts and objects.

- Select **No** for **Migrate Account**.
- Select **All** for **Migrate Object**.

Step 11 Click **Next**. 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**.
- If all check items are successful, click **Next**.

Step 12 Click **Submit**.

Return to the **Online Migration Management** page and check the migration task status.

It takes several minutes to complete.

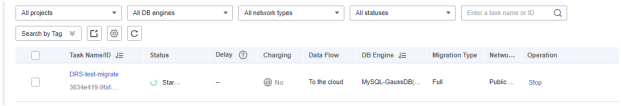

If the status changes to **Completed**, the migration task has been created.

NOTE

- Currently, MySQL to GaussDB(for MySQL) migration supports two modes: full migration and full+incremental migration.
- If you create a full migration task, the task automatically stops after the full data is migrated to the destination.
- If you create a full+incremental migration task, a full migration is executed first. After the full migration is complete, an incremental migration starts.
- During the incremental migration, data is continuously migrated so the task will not automatically stop.

----End

2.1.8 Checking Migration Results

You can use either of the following methods to check the migration results:

1. DRS compares migration objects, users, and data and provide comparison results. For details, see **Checking the Migration Results on the DRS Console**.

 Copyright © Huawei Technologies Co., Ltd.

2. Log in to the destination side to check whether the databases, tables, and data are migrated. Confirm the data migration status. For details, see **Checking the Migration Results on the GaussDB Console**.

Checking the Migration Results on the DRS Console

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select region EU-Dublin.

Step 3 Under the service list, choose **Databases** > **Data Replication Service**.

Step 4 Click the DRS instance name.

Step 5 Click **Migration Comparison** and select **Object-Level Comparison** to check whether database objects are missing.

Step 6 Choose **Data-Level Comparison** and check whether the number of rows of migrated objects is consistent.

Step 7 Click **Account-Level Comparison** and check whether the accounts and permissions of the source and destination instances are the same.

----End

Checking the Migration Results on the GaussDB Console

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select region EU-Dublin.

Step 3 Under the service list, choose **Databases** > **GaussDB**.

Step 4 Select **GaussDB(for MySQL)**, locate the destination instance, and click **Log In** in the **Operation** column.

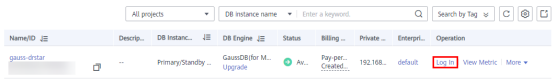

Step 5 In the dialog box that is displayed, enter the password and click **Test Connection**.

Step 6 After the connection is successful, click **Log In**.

Step 7 Check whether the destination databases and tables are the same as the source instance. Check whether migration is complete.

----End

2.2 From Other Cloud MongoDB to DDS

DRS helps you migrate MongoDB databases from other cloud platforms to DDS on the current cloud. With DRS, you can migrate databases online with zero downtime and your services and databases can remain operational during migration.

 Copyright © Huawei Technologies Co., Ltd.

This section describes how to use DRS to migrate MongoDB databases from another cloud to DDS on the current cloud. Migration scenarios include:

- Migrating MongoDB databases from another cloud to DDS on the current cloud.
- Migrating self-built MongoDB databases from servers on another cloud to DDS on the current cloud.

Diagram

Figure 2-3 Migrating MongoDB databases from other clouds

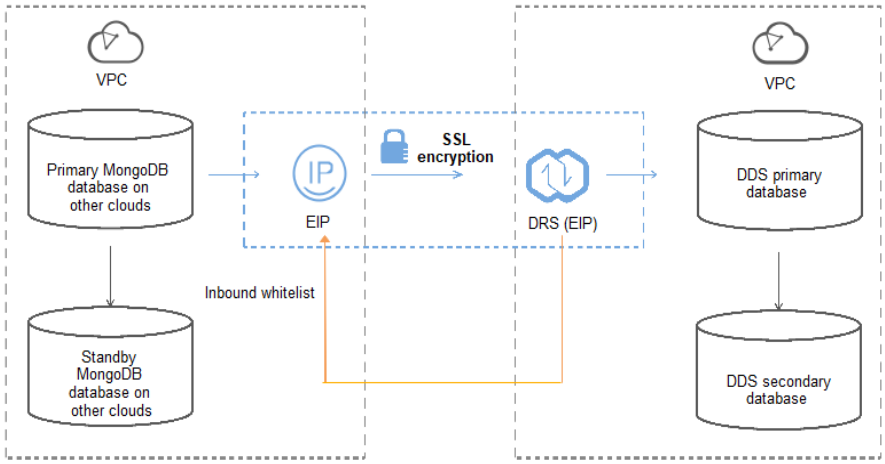

Figure 2-4 Migrating MongoDB databases from other cloud servers

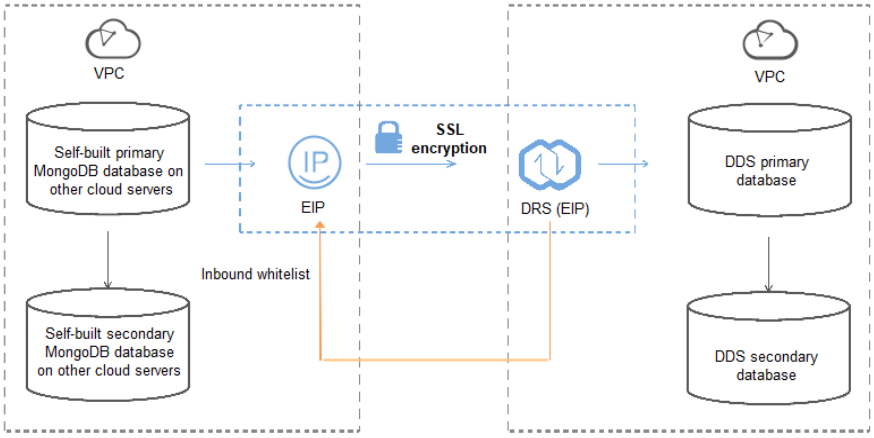

 Copyright © Huawei Technologies Co., Ltd.

Migration Process

Figure 2-5 Flowchart

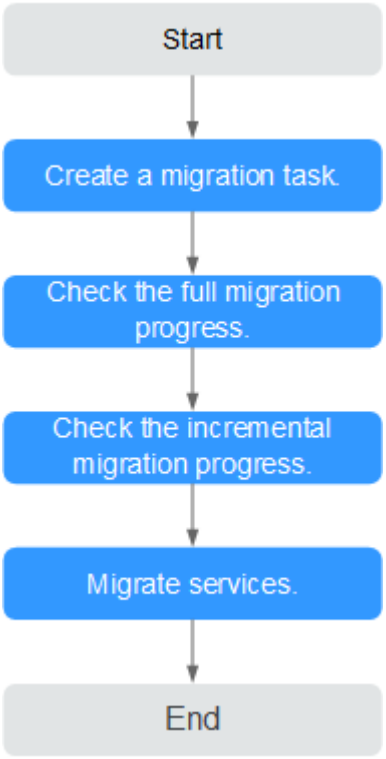

Migration Suggestions (Important)

- Database migration is closely impacted by a wide range of environmental and operational factors. To ensure the migration goes smoothly, perform a test run before the actual migration to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.
- It is strongly recommended that you 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.

Notes on Migration (Important)

NOTICE

Before creating a migration task, read the migration notes carefully.

For details, see **precautions** on using specific migration tasks in *Data Replication Service Real-Time Migration*.

 Copyright © Huawei Technologies Co., Ltd.

Preparations

1. Permissions

   Table 2-3 lists the permissions required for the source and destination databases when migrating a MongoDB database from another cloud to DDS on the current cloud.

   Table 2-3 Migration permissions

Database	Full Migration Permission	Full+Incremental Migration Permission
Source	• Replica set: The source database user must have the read permission for the database to be migrated. • Single node: The source database user must have the read permission for the database to be migrated. • Cluster: The source database user must have the read permission for the databases to be migrated and the config database. • To migrate accounts and roles of the source database, the source database user must have the read permission for the **system.users** and **system.roles** system tables of the admin database.	• Replica set: The source database user must have the read permission for the databases to be migrated and the local database. • Single node: The source database user must have the read permission for the databases to be migrated and the local database. • Cluster: The source mongos node user must have the readAnyDatabase permission for the databases to be migrated and 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 database user must have the read permission for the **system.users** and **system.roles** system tables of the admin database.

 Copyright © Huawei Technologies Co., Ltd.

Database	Full Migration Permission	Full+Incremental Migration Permission
Destination	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 migration account must have the read permission for the config database.	

- Source database permissions:

  The source MongoDB database user must have all the required permissions listed in **Table 2-3**. If the permissions are insufficient, create a user that has all of the permissions on the source database.

  Data of the following third-party cloud vendors can be migrated: Alibaba Cloud and Tencent Cloud (Tencent Cloud 3.2 is not supported).

  For example, if you want to migrate an Alibaba Cloud MongoDB database to a DDS database, you can use the initial account of the source database.
- Destination database permissions:

  If the destination database is a DDS database, the initial account can be used.

2. Network settings

   Enable public accessibility for the source database.

   - Source database network settings:

     Any source database MongoDB instances will need to be accessible from the Internet.
   - Destination database network settings: No settings are required.

3. Security rules

   - Source database security group settings:

     The replication instance needs to be able to access the source MongoDB instance. That means that the EIP of the replication instance must be on the whitelist of the source MongoDB instance.

     Before configuring the network whitelist, you need to obtain the EIP of the replication instance.

     - After creating a replication instance on the DRS console, you can find the EIP on the **Configure Source and Destination Databases** page as shown in **Figure 2-6**.

       Figure 2-6 EIP of the replication instance

       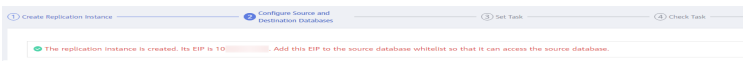

     You can also add 0.0.0.0/0 to the source database whitelist to allow any IP address to access the source database but this action may result in security risks.

 Copyright © Huawei Technologies Co., Ltd.

If you do take this step, then once the migration is complete, you should delete this item from the whitelist or your system will insecure.

- Destination database security group settings:

  By default, the destination database and the DRS replication instance are in the same VPC and can communicate with each other. No further configuration is required.

4. Other

   You need to export the user information of the MongoDB database first and manually add it to the destination DDS DB instance because the user information will not be migrated.

Migration Procedure

Step 1 Create a migration task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.
2. On the **Online Migration Management** page, click **Create Migration Task**.
3. On the **Replication Instance Information** page, configure the task details, description, and replication instance details and click **Next**.

Figure 2-7 Replication instance information

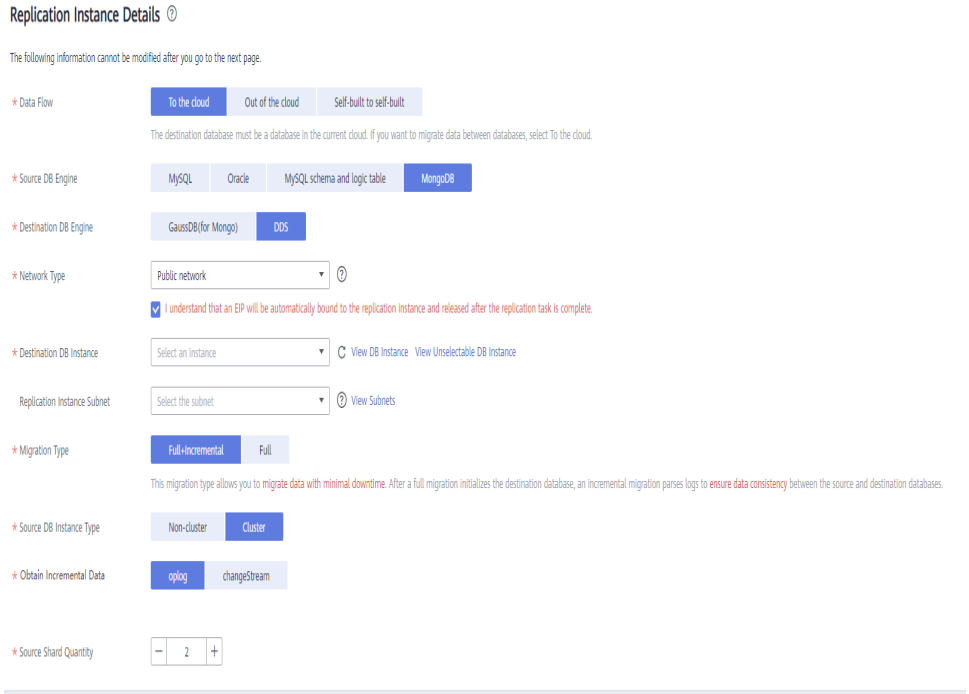

Table 2-4 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.

 Copyright © Huawei Technologies Co., Ltd.

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 consists of a maximum of 256 characters and cannot contain the following special characters: =<>&'\"

Table 2-5 Replication instance settings

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	Select **MongoDB**.
Destination DB Engine	Select **DDS**.
Network Type	Select **Public network**.
Destination DB Instance	The DDS DB instance you purchased.
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. The objects include collections and indexes. – **Full+Incremental** The full+incremental migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration parses logs to ensure data consistency between the source and destination databases.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
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 mode, you must disable the balancer of the source instance. When testing the connection, you need to enter the connection information of each shard node of the source instance. – changeStream: This method is recommended. For MongoDB 4.0 and later, DRS connects to mongos nodes of the source instance to extract data. If you select this method, you must enable the WiredTiger storage engine of the source instance. **NOTE** Only whitelisted users can use **changeStream**. To use this function, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets** > **Create Service Ticket** to submit a service ticket.
Source Shard Quantity	If **Source DB Instance Type** is set to **Cluster** and **Obtain Incremental Data** is set to **oplog**, enter the number of source shard nodes. The default minimum number of source DB instances is 2 and the maximum number is 32. You can set this parameter based on the number of source database shards.
Tags	This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.

4. 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**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-8 Source database information

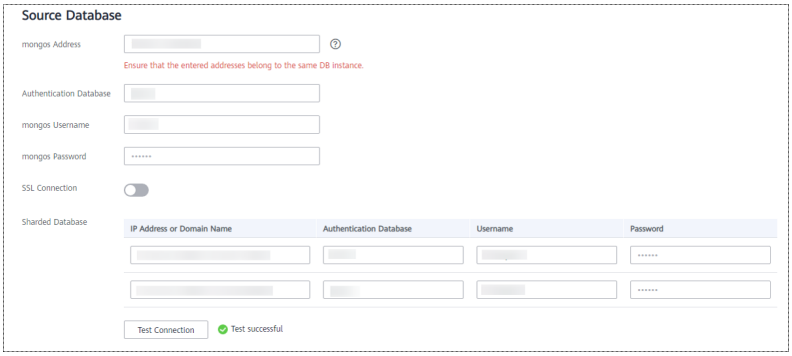

Table 2-6 Source database settings

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	A username for the source database.
mongos Password	The password for the source database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.
Sharded Database	Enter the information about the sharded databases in the source database.

- Destination database configuration

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-9 Destination database information

Table 2-7 Destination database settings

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

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

Figure 2-10 Migration object

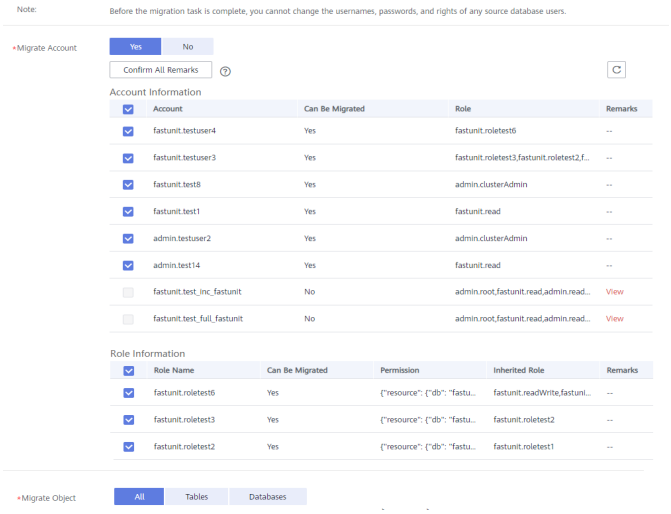

 Copyright © Huawei Technologies Co., Ltd.

Table 2-8 Migration object

Parameter	Description
Migrate Account	There are accounts that can be migrated completely 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 choose to migrate accounts, see **Migrating Accounts** in *Data Replication Service User Guide* to migrate database users and roles. – **No** During the migration, accounts and roles are not migrated.
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 database to be migrated may have dependencies on other objects that are not migrated. To ensure a successful migration, you are advised to 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.

6. 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*.
   – If all check items are successful, click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-11 Task Check

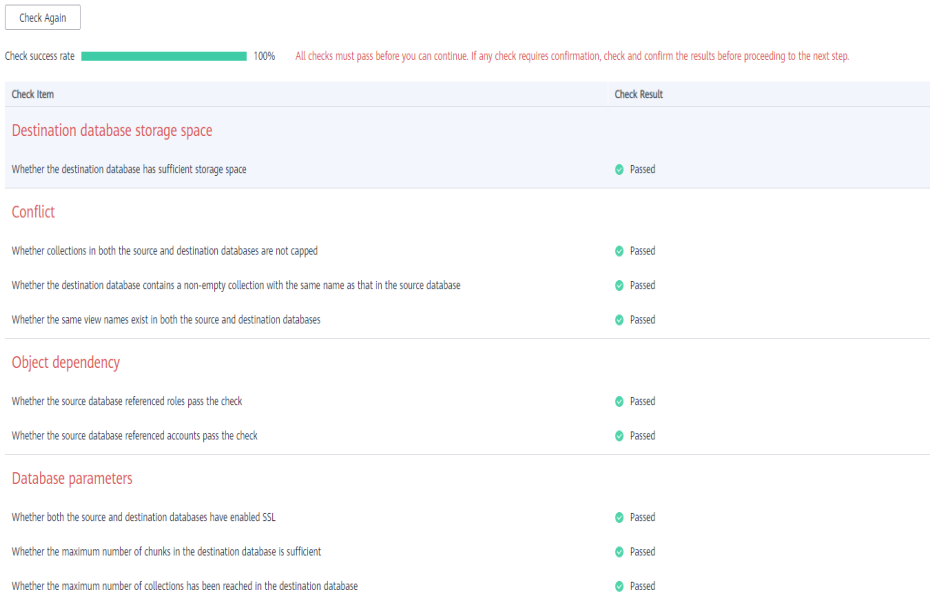

NOTE

You can proceed to the next step only when all check items are successful. If any alarms are generated, view and confirm the alarm details first before proceeding to the 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

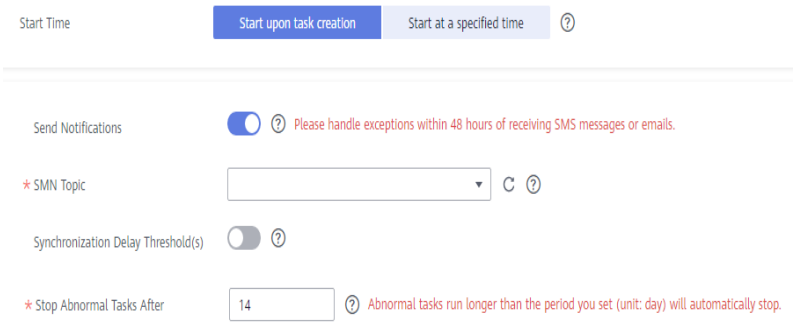

 Copyright © Huawei Technologies Co., Ltd.

Table 2-9 Task startup settings

Parameter	Description
Start Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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 Notification** 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.

8. After the task is submitted, go back to the **Online Migration Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full migration

 Copyright © Huawei Technologies Co., Ltd.

  - Viewing the migration progress: Click the target full migration task, and on the **Migration Progress** tab, you can see the migration progress of the structure, data, indexes, and migration objects. When the progress reaches 100%, the migration is complete.
  - Viewing migration details: 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. Currently, this function is available only to whitelisted users. You can submit a service ticket to apply for this function.
- Incremental Migration Permission
  - Viewing the synchronization delay: After the full migration is complete, an incremental migration starts. On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Progress** to view the synchronization delay of the incremental migration. If the synchronization delay is 0s, the destination database is being synchronized with the source database in real time. You can also view the data consistency on the **Migration Comparison** tab.

    Figure 2-13 Viewing the synchronization delay

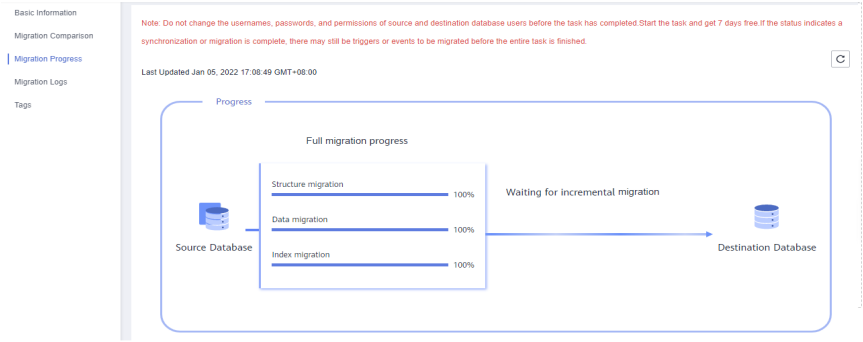

  - Viewing the migration results: On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Comparison** and perform a migration comparison in accordance with the comparison process, which should help you determine an appropriate time for migration to minimize service downtime.

    Figure 2-14 Database comparison process

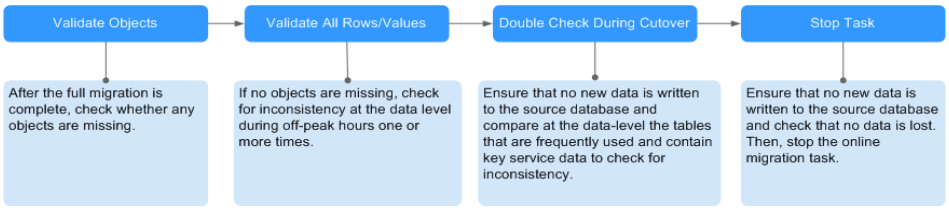

    For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. To obtain accurate

 Copyright © Huawei Technologies Co., Ltd.

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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.
2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.
   ```
   db.currentOp()
   ```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS replication instance. If no additional session executes SQL statements, the service has been stopped.
3. On the **Migration Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the migration is complete.

Step 4 Stop or delete the migration task.

1. Stopping the migration task. After databases and services are migrated to the destination database, to prevent operations on the source database from being synchronized to the destination database to overwrite data, you can stop the migration task. This operation only deletes the replication instance, and the migration task is still displayed in the task list. You can view or delete the task. After the task is stopped, DRS will not charge you anymore.
2. Delete the migration task. After the migration task is complete, you can delete it. After the migration task is deleted, it will no longer be displayed in the task list.

----End

2.3 From MongoDB on ECS to DDS

DRS helps you migrate data from MongoDB databases on ECSs to DDS instances on the current cloud. With DRS, you can migrate databases online with zero downtime and your services and databases can remain operational during migration.

This section describes how to use DRS to migrate data from an ECS database to a DDS instance on the current cloud. The following network scenarios are supported:

 Copyright © Huawei Technologies Co., Ltd.

- Source and destination databases are in the same VPC.
- Source and destination databases are in different VPCs.

Diagram

Figure 2-15 Source and destination databases in the same VPC

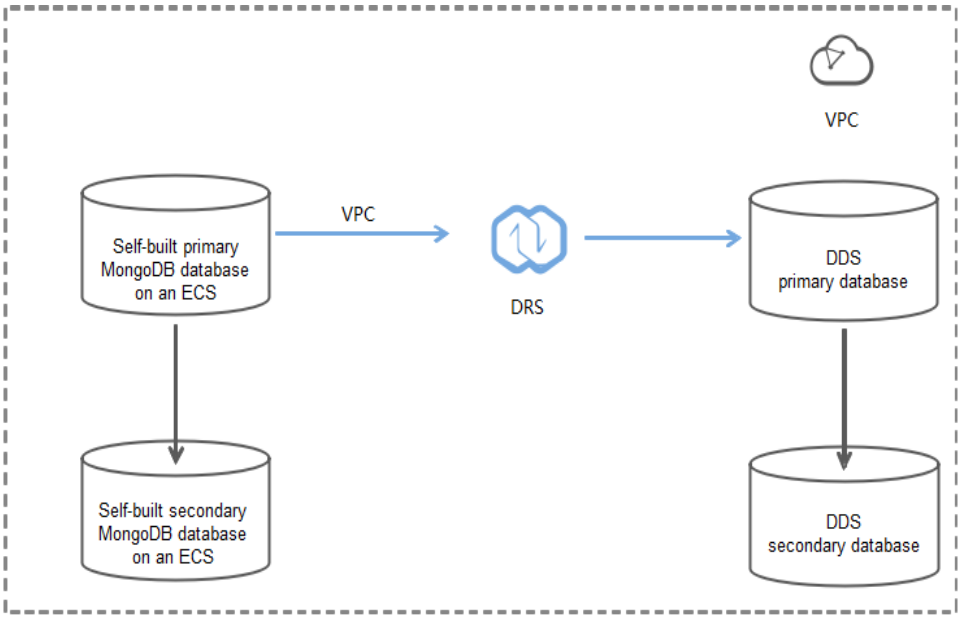

Figure 2-16 Source and destination databases in the same region and different VPCs

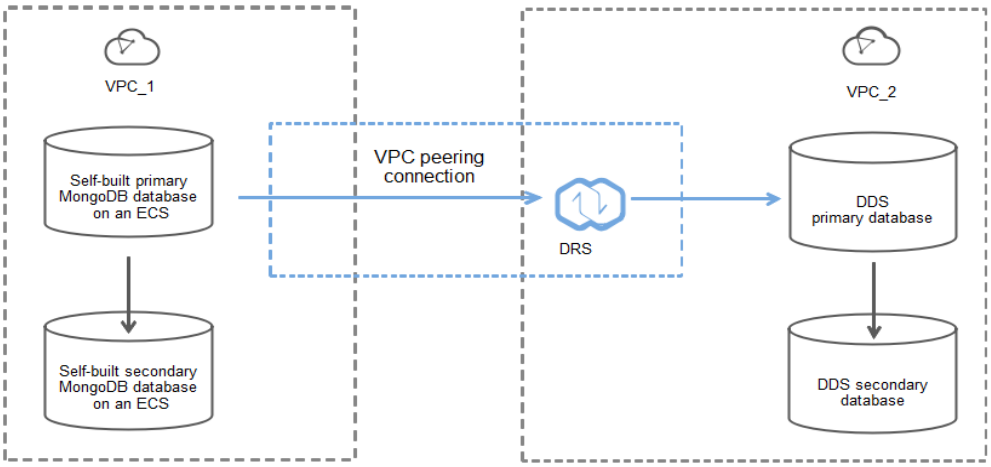

 Copyright © Huawei Technologies Co., Ltd.

Migration Process

Figure 2-17 Flowchart

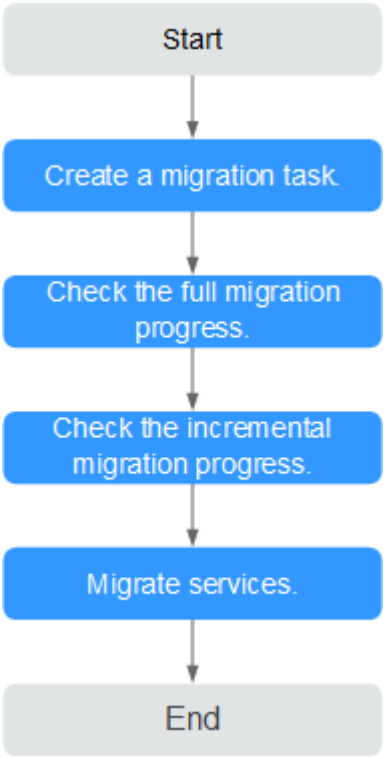

Migration Suggestions (Important)

- Database migration is closely impacted by a wide range of environmental and operational factors. To ensure the migration goes smoothly, perform a test run before the actual migration to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.
- It is strongly recommended that you 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.

Notes on Migration (Important)

NOTICE

Before creating a migration task, read the migration notes carefully.

For details, see **precautions** on using specific migration tasks in *Data Replication Service Real-Time Migration*.

 Copyright © Huawei Technologies Co., Ltd.

Preparations

1. Permissions:

   Table 2-10 lists the permissions required for the source and destination databases when migrating data from a MongoDB database on an ECS to DDS on the current cloud.

   Table 2-10 Migration permissions

Database	Full Migration Permission	Full+Incremental Migration Permission
Source	• Replica set: The source database user must have the read permission for the database to be migrated. • Single node: The source database user must have the read permission for the database to be migrated. • Cluster: The source database user must have the read permission for the databases to be migrated and the config database. • To migrate accounts and roles of the source database, the source database user must have the read permission for the **system.users** and **system.roles** system tables of the admin database.	• Replica set: The source database user must have the read permission for the databases to be migrated and the local database. • Single node: The source database user must have the read permission for the databases to be migrated and the local database. • Cluster: The source mongos node user must have the readAnyDatabase permission for the databases to be migrated and 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 database user must have the read permission for the **system.users** and **system.roles** system tables of the admin database.

 Copyright © Huawei Technologies Co., Ltd.

Database	Full Migration Permission	Full+Incremental Migration Permission
Destination	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 migration account must have the read permission for the config database.	

- Source database permissions:

  The source MongoDB database user must have all the required permissions listed in **Table 2-10**. If the permissions are insufficient, create a user that has all of the permissions on the source database.

- Destination database permissions:

  The initial account of the DDS instance has the required permissions.

2. Network settings
   - The source database and destination DDS DB instance must be in the same region.
   - The source database and destination DDS DB instance can be either in the same VPC or different VPCs.
     - If the source and destination databases are in different VPCs, the subnets of the source and destination databases are required to be in different CIDR blocks. You need to create a VPC peering connection between the two VPCs. For details, see **VPC Peering Connection Creation Procedure**.
     - If the source and destination databases are in the same VPC, the networks are interconnected by default.
3. Security rules
   - In the same VPC, the network is connected by default. You do not need to set a security group.
   - In different VPCs, establish a VPC peering connection between the two VPCs. You do not need to set a security group.
4. Other

   You need to export the user information of the MongoDB database first and manually add it to the destination DDS DB instance because the user information will not be migrated.

Migration Procedure

Step 1 Create a migration task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.
2. On the **Online Migration Management** page, click **Create Migration Task**.

 Copyright © Huawei Technologies Co., Ltd.

3. On the **Create Replication Instance** page, configure the task details, recipient, and replication instance and click **Next**.

Figure 2-18 Replication instance information

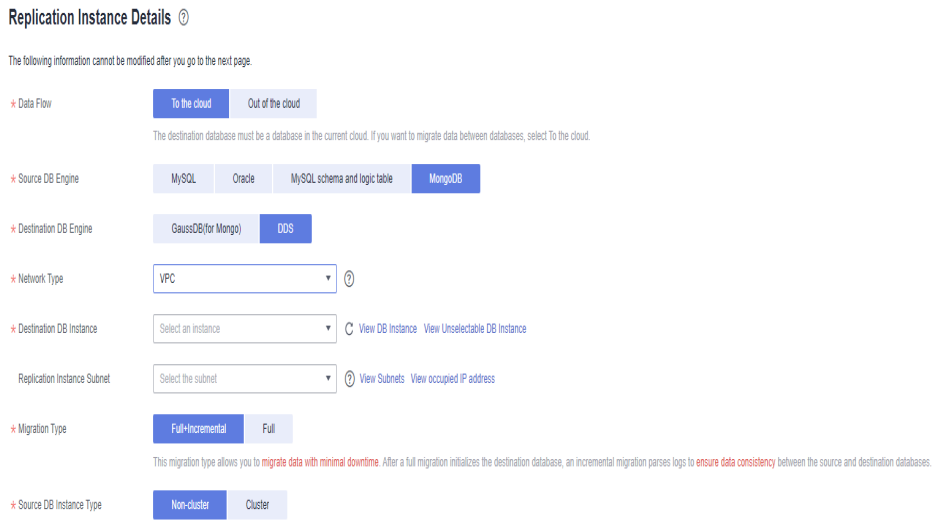

Table 2-11 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.
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: =<>&'\"

Table 2-12 Replication instance information

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	**MongoDB database**
Destination DB Engine	**DDS**
Network Type	**VPC**
Destination DB Instance	The DDS DB instance you purchased.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Migration Type	Select **Full+Incremental** as an example: – **Full**: This migration type is suitable for scenarios where a 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 perform a full migration, you are advised to stop 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 the **Full+Incremental** migration type, data generated during the full migration will be synchronized to the destination database with zero downtime, ensuring that both the source and destination databases remain accessible.
Source DB Instance Type	If you select **Full+Incremental** for **Migration Type**, set this parameter based on the source database. **Non-cluster** is selected as an example. – 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**.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
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 mode, you must disable the balancer of the source instance. When testing the connection, you need to enter the connection information of each shard node of the source instance. – changeStream: This method is recommended. For MongoDB 4.0 and later, DRS connects to mongos nodes of the source instance to extract data. If you select this method, you must enable the WiredTiger storage engine of the source instance. **NOTE** Only whitelisted users can use **changeStream**. To use this function, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets** > **Create Service Ticket** to submit a service ticket.
Source Shard Quantity	If **Source DB Instance Type** is set to **Cluster** and **Obtain Incremental Data** is set to **oplog**, enter the number of source shard nodes. The default minimum number of source DB instances is 2 and the maximum number is 32. You can set this parameter based on the number of source database shards.

4. 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**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-19 Source and destination database details

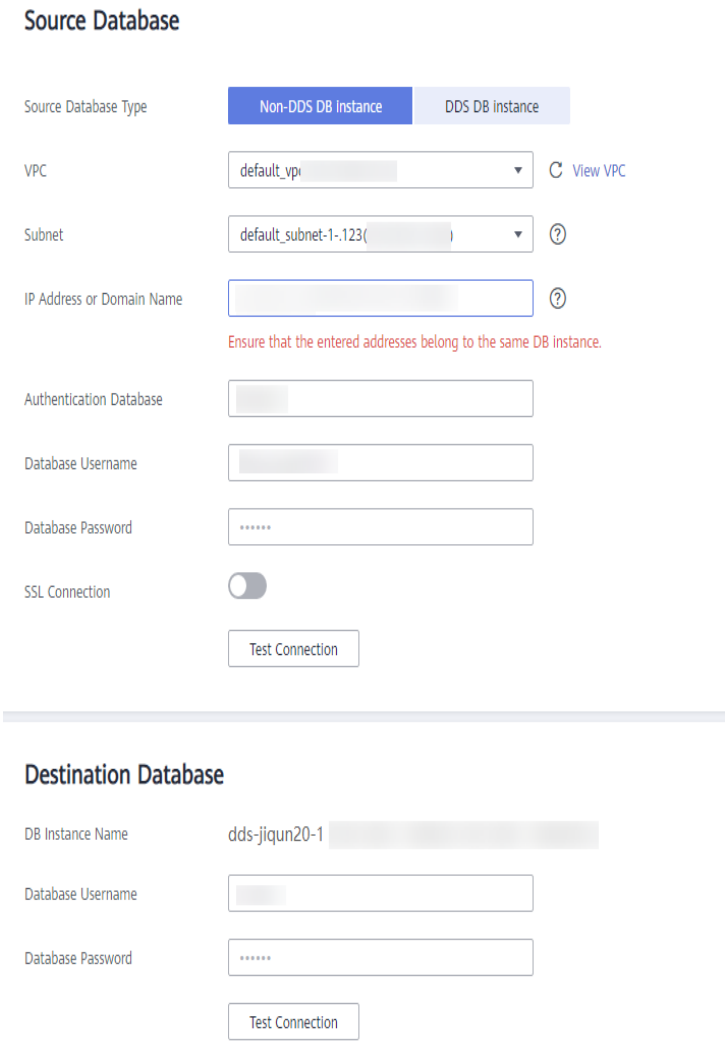

Table 2-13 Source database information

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. For details on how to create a VPC, see **Creating a VPC**.
Subnet	A subnet provides dedicated network resources that are logically 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. For details on how to create a VPC, see the **Creating a VPC** section in the *Virtual Private Cloud User Guide*.

 Copyright © Huawei Technologies Co., Ltd.

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	A username for the source database.
Database Password	The password for the database username.
SSL Connection	To improve data security during the migration, you are advised to enable SSL to encrypt migration links and upload a CA certificate.

Table 2-14 Destination database information

Parameter	Description
DB Instance Name	The DDS DB instance you have selected during the migration task creation is displayed by default and cannot be changed.
Database Username	The username for accessing the destination DDS DB instance.
Database Password	The password for the database username.

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

Figure 2-20 Migration object

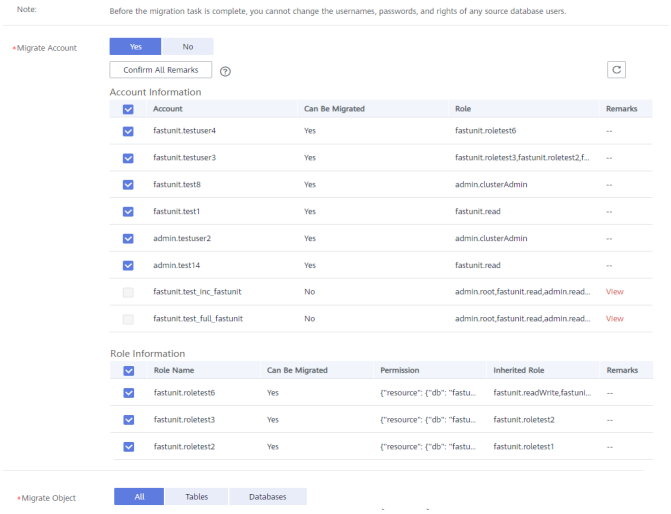

 Copyright © Huawei Technologies Co., Ltd.

Table 2-15 Migration object

Parameter	Description
Migrate Account	There are accounts that can be migrated completely 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 choose to migrate accounts, see **Migrating Accounts** in *Data Replication Service User Guide* to migrate database users and roles. – **No** During the migration, accounts and roles are not migrated.
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 database to be migrated may have dependencies on other objects that are not migrated. To ensure a successful migration, you are advised to 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.

6. 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*.
   - If all check items are successful, click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-21 Task Check

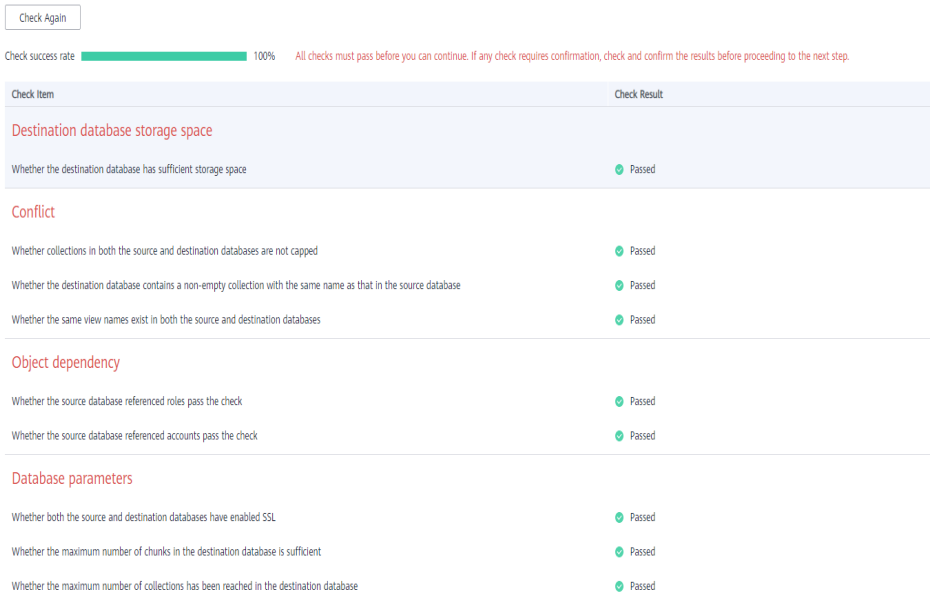

NOTE

You can proceed to the next step only when all check items are successful. If any alarms are generated, view and confirm the alarm details first before proceeding to the 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 2-22 Task startup settings

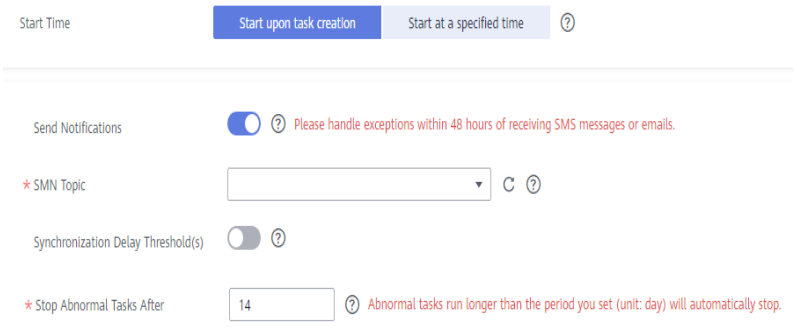

 Copyright © Huawei Technologies Co., Ltd.

Table 2-16 Task startup settings

Parameter	Description
Start Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. The **Start at a specified time** option is recommended. **NOTE** The migration task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.
Send Notifications	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 Notification** 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.

8. After the task is submitted, go back to the **Online Migration Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full migration

 Copyright © Huawei Technologies Co., Ltd.

  - Viewing the migration progress: Click the target full migration task, and on the **Migration Progress** tab, you can see the migration progress of the structure, data, indexes, and migration objects. When the progress reaches 100%, the migration is complete.
  - Viewing migration details: 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. Currently, this function is available only to whitelisted users. You can submit a service ticket to apply for this function.
- Incremental Migration Permission
  - Viewing the synchronization delay: After the full migration is complete, an incremental migration starts. On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Progress** to view the synchronization delay of the incremental migration. If the synchronization delay is 0s, the destination database is being synchronized with the source database in real time. You can also view the data consistency on the **Migration Comparison** tab.

    Figure 2-23 Viewing the synchronization delay

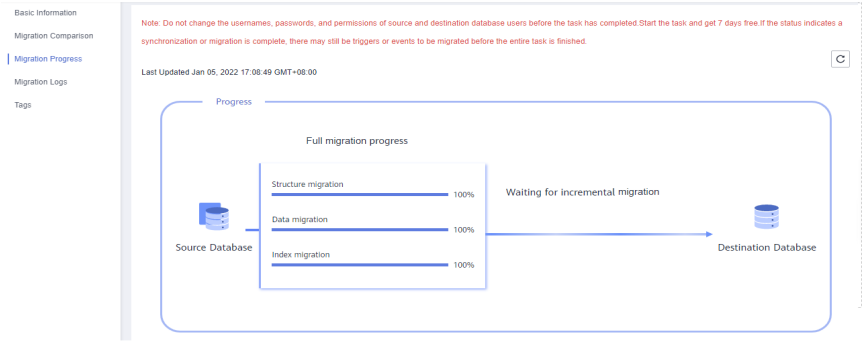

  - Viewing the migration results: On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Comparison** and perform a migration comparison in accordance with the comparison process, which should help you determine an appropriate time for migration to minimize service downtime.

    Figure 2-24 Database comparison process

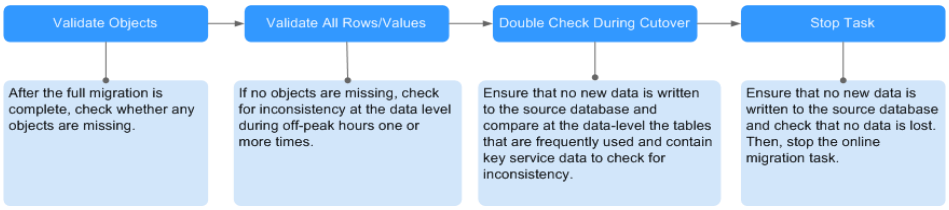

    For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. To obtain accurate

 Copyright © Huawei Technologies Co., Ltd.

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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.
2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.
   ```
   db.currentOp()
   ```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS replication instance. If no additional session executes SQL statements, the service has been stopped.
3. On the **Migration Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the migration is complete.

Step 4 Stop or delete the migration task.

1. Stopping the migration task. After databases and services are migrated to the destination database, to prevent operations on the source database from being synchronized to the destination database to overwrite data, you can stop the migration task. This operation only deletes the replication instance, and the migration task is still displayed in the task list. You can view or delete the task. After the task is stopped, DRS will not charge you anymore.
2. Delete the migration task. After the migration task is complete, you can delete it. After the migration task is deleted, it will no longer be displayed in the task list.

----End

2.4 From On-Premises MySQL to RDS MySQL

DRS supports data migration from on-premises MySQL databases to RDS MySQL instances on the current cloud. With DRS, you can migrate databases online with zero downtime and your services and databases can remain operational during migration.

This section describes how to use DRS to migrate data from an on-premises MySQL database to an RDS MySQL instance on the current cloud. The following network types are supported:

 Copyright © Huawei Technologies Co., Ltd.

- Virtual Private Network (VPN)
- Public network

Diagram

Figure 2-25 VPN

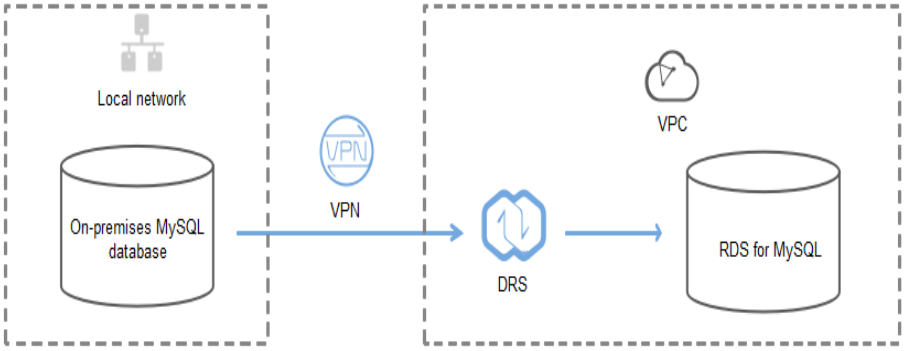

Figure 2-26 Public network+SSL connection

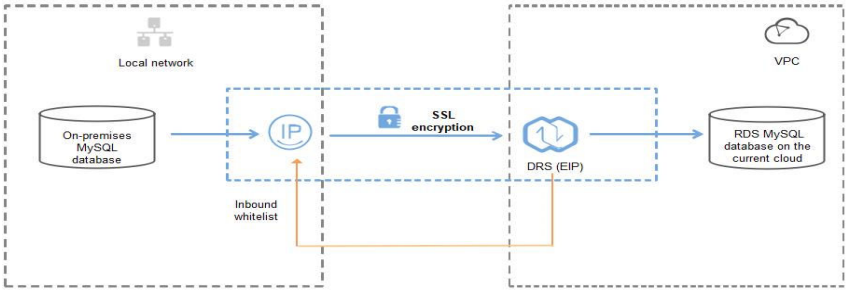

 Copyright © Huawei Technologies Co., Ltd.

Migration Process

Figure 2-27 Flowchart

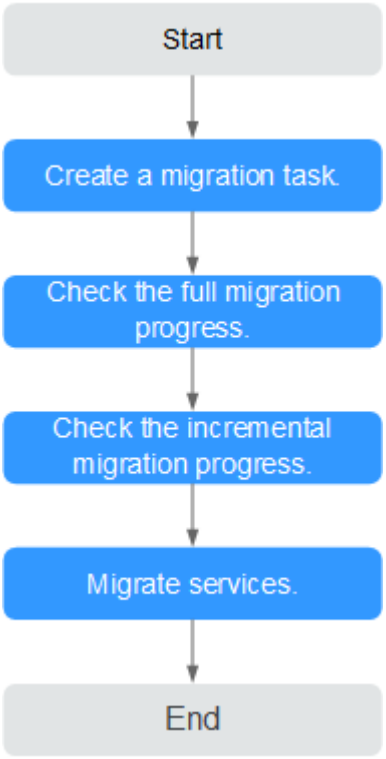

Migration Suggestions (Important)

- Database migration is closely impacted by a wide range of environmental and operational factors. To ensure the migration goes smoothly, perform a test run before the actual migration to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.
- It is strongly recommended that you 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.

Notes on Migration (Important)

NOTICE

Before creating a migration task, read the migration notes carefully.

For details, see **precautions** on using specific migration tasks in *Data Replication Service Real-Time Migration*.

 Copyright © Huawei Technologies Co., Ltd.

Preparations

1. Permissions

   Table 2-17 lists the permissions required for the source and destination databases when migrating data from on-premises MySQL databases to the RDS MySQL instances on the current cloud.

   Table 2-17 Required permissions

Database	Full Migration	Full+Incremental Migration
Source	SELECT, SHOW VIEW, and EVENT	SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT
Destination	SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, and WITH GRANT OPTION.	SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, and WITH GRANT OPTION.

   - Source database permissions:

     The source database user must have all the required permissions listed in Table 2-17. If the permissions are insufficient, create a user that has all of the permissions on the source database.
   - Destination database permissions:

     If the destination database is an RDS MySQL database on the current cloud, the initial account can be used.

2. Network settings
   - Source database network settings:

     You can migrate data from on-premises MySQL databases to an RDS MySQL instance on the current cloud through a VPN or public network. Enable public accessibility or establish a VPN for the on-premises MySQL databases based on your service requirements. You are advised to migrate data through a public network, which is more convenient and cost-effective.
   - Destination database network settings:
     - If the source database attempts to access the destination database through a VPN, you need to enable the VPN service first so that the source database can communicate with the destination RDS MySQL.
     - If the source database attempts to access the destination database through a public network, you do not need to configure the destination database.

 Copyright © Huawei Technologies Co., Ltd.

3. Security rules
   a. Source database security group settings:
      - The replication instance needs to be able to access the source DB. That means that the EIP of the replication instance must be on the whitelist of the source MySQL DB instance. Before configuring the network whitelist, you need to obtain the EIP of the replication instance.

        After creating a replication instance on the DRS console, you can find the EIP on the **Configure Source and Destination Databases** page.
      - If the migration is performed over a VPN network, add the private IP address of the DRS replication instance to the network whitelist of the source MySQL database to enable the source MySQL database to communicate with the current cloud. The IP address on the **Configure Source and Destination Databases** page is the private IP address of the replication instance.

      If you do take this step, then once the migration is complete, you should delete this item from the whitelist or your system will insecure.
   b. Destination database security group settings:

      By default, the destination database and the DRS replication instance are in the same VPC and can communicate with each other. No further configuration is required.
4. Other

   DRS supports migration of some parameters that are closely related to services and performance. For details about these parameters, see **Parameters for Comparison**. If you need to migrate other parameters, configure them manually based on service requirements.

Migration Procedure

The following describes how to use DRS to migrate data from an on-premises MySQL database to an RDS MySQL instance on the current cloud over a public network.

Step 1 Create a migration task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.
2. On the **Online Migration Management** page, click **Create Migration Task**.
3. On the **Create Replication Instance** page, configure the task details and the replication instance, and click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-28 Replication instance information

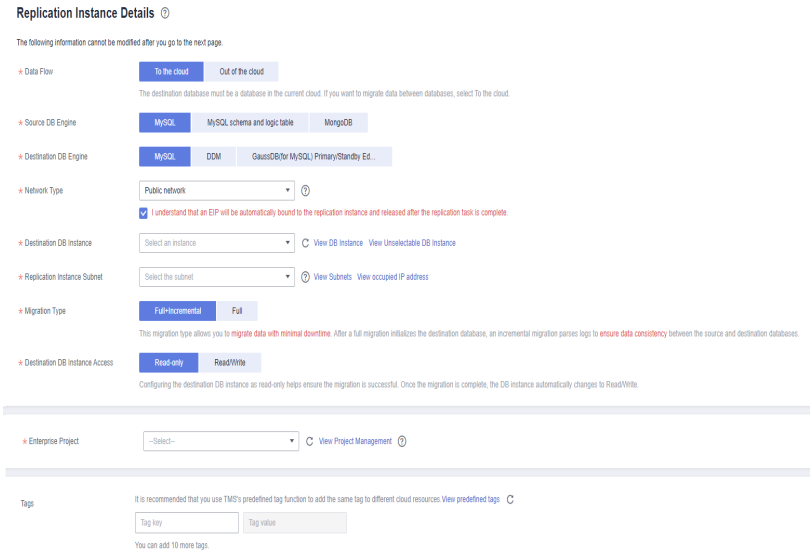

Table 2-18 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.
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: =<>&'\"

Table 2-19 Replication instance settings

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	**MySQL**
Destination DB Engine	**MySQL**
Network Type	Select **Public network**. Enabling SSL is recommended. It may slow down the migration by 20% to 30% but it ensures data security.
Destination DB Instance	The MySQL DB instance you purchased.

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Destination Database Access	You can select **Read-only** or **Read/Write**. – **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 migration, the destination database can be queried or modified. Data 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.
Migration Type	Select **Full+Incremental** as an example. – **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 the **Full+Incremental** migration type, data generated during the full migration will be synchronized to the destination database with zero downtime, ensuring that the source database remain accessible.

4. 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**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-29 Source and destination database details

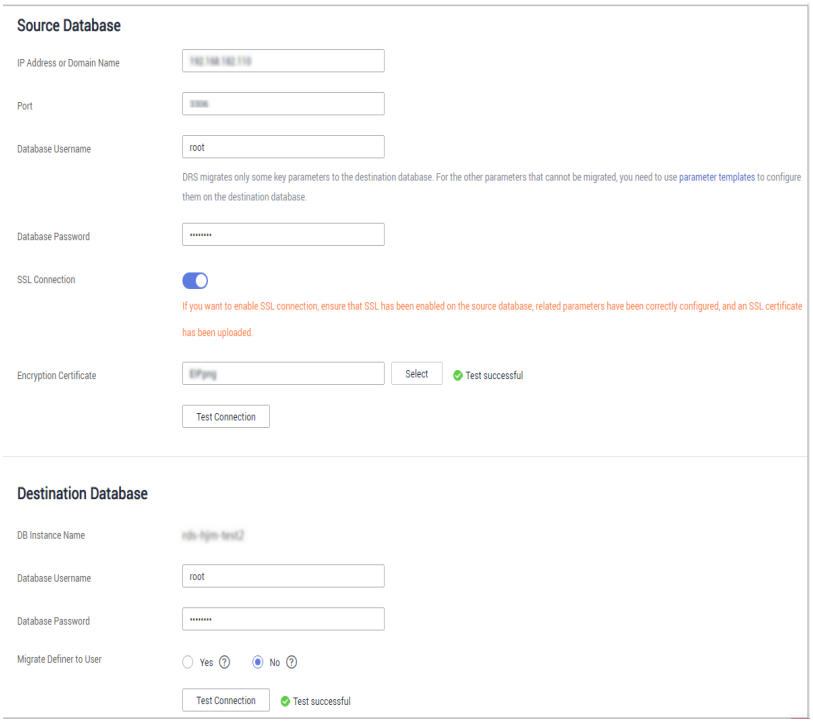

Table 2-20 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	Enter an integer ranging from 1 to 65535, which indicates the port number of the source database.
Database Username	A username for the source database.
Database Password	The password for the database username.
SSL Connection	To improve data security during a migration on a public network, you are advised to enable SSL to encrypt migration links and upload a CA certificate.

Table 2-21 Destination database settings

Parameter	Description
DB Instance Name	The RDS MySQL instance you selected during the replication instance creation is displayed by default and cannot be changed.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Database Username	The username for accessing the destination MySQL instance.
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 on authorization, see **How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration?** in *Data Replication Service FAQs*. – **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.

5. On the **Set Task** page, set the flow control mode and select migration accounts and objects.

Figure 2-30 Migration object

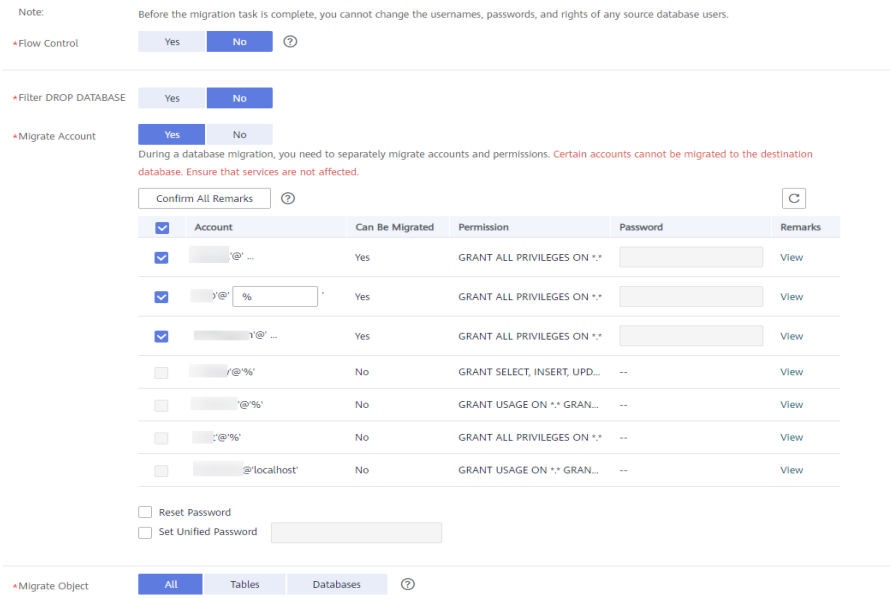

 Copyright © Huawei Technologies Co., Ltd.

Table 2-22 Migration types and objects-public network

Parameter	Description
Flow Control	You can choose whether to control the flow. – 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. – If the migration speed is not limited, 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. In the task list on the **Online Migration Management** page, locate the target task and choose **More** > **Speed** or **Speed** in the **Operation** column.
Filter DROP DATABASE	During synchronization, executing DDL operations on the source database may affect the data synchronization performance to some extent. To reduce data synchronization risks, DRS allows you to filter out DDL operations. The database deletion operation can be filtered out by default. – If you select **Yes**, the database deletion operation performed on the source database is not synchronized during data synchronization. – If you select **No**, related operations are synchronized to the destination database during data synchronization. **NOTE** Only the database deletion operation can be filtered.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
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 your service requirements. – **Yes** If you choose to migrate accounts, see **Migrating Accounts** in *Data Replication Service User Guide* to migrate database users, permissions, and passwords. – **No** During the migration, accounts and permissions are not migrated.
Migrate Object	All database objects can be migrated. After the objects are migrated to the destination DB instance, the object names remain the same as those in the source database and cannot be modified. You can migrate all objects or specified objects based on your service requirements. – **All**: All objects in the source database are migrated to the destination database. – **Self-defined**: Only self-defined objects are migrated to the destination 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 database to be migrated may have dependencies on other objects that are not migrated. To ensure a successful migration, you are advised to migrate all of the databases.

6. Click **Next**. 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**.
   - If all check items are successful, click **Next**.
7. Compare the source and target 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.
   - As this process is optional, you can click **Next** if you wish to skip this step.
   - Compare common parameters:

     If the parameter values in the list shown are inconsistent, you can click **Save Change** to change the destination database values to match those of the source database.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-31 Modifying common parameters

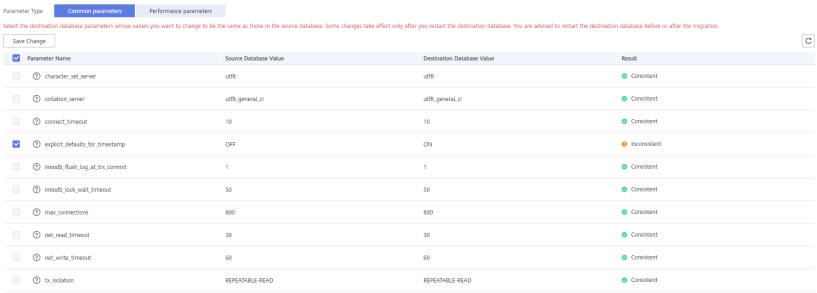

Performance parameter values in both the source and destination can be the same or different.

- There is a value that is consistent, but you still want to change it in the destination, locate the parameter, enter the value in the **Change To** column, and click **Save Change** in the upper left corner.
- If you want to change a destination database value to match the source same:

  1) Click **Use Source Database Value**.

     The system automatically updates the destination database value to match the source.

     Figure 2-32 One-click modification

     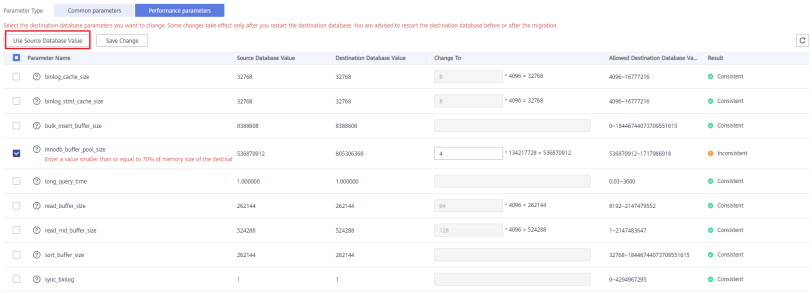

     NOTE

     You can also manually enter parameter values.

  2) 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 comparison results are automatically updated.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-33 Performance parameters

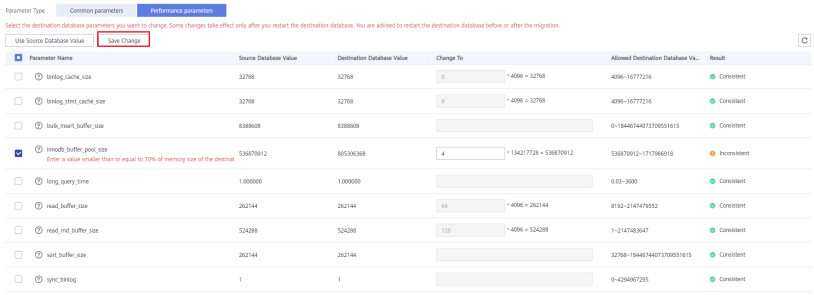

Some parameters in the destination database require a restart before the changes can take effect. The system will display these as being inconsistent. You will need to restart the destination database after either before the migration starts or after it has completed. To minimize the impact of this restart 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 parameter comparison, see **Parameters for Comparison** in the *Data Replication Service User Guide*.

3) Click **Next**.

8. 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

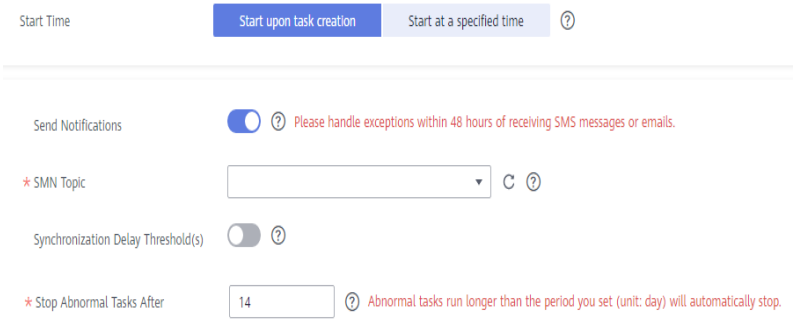

Table 2-23 Task startup settings

Parameter	Description
Start Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. The **Start at a specified time** option is recommended. **NOTE** The migration task may affect the performance of the source and destination databases. You are advised to start the task in off-peak hours and reserve two to three days for data verification.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
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 Notification** 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.

9. After the task is submitted, go back to the **Online Migration Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full migration
  - Viewing the migration progress: Click the target full migration task, and on the **Migration Progress** tab, you can see the migration progress of the structure, data, indexes, and migration objects. When the progress reaches 100%, the migration is complete.
  - Viewing migration details: 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. Currently, this

 Copyright © Huawei Technologies Co., Ltd.

function is available only to whitelisted users. You can submit a service ticket to apply for this function.

- Incremental Migration Permission
  - Viewing the synchronization delay: After the full migration is complete, an incremental migration starts. On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Progress** to view the synchronization delay of the incremental migration. If the synchronization delay is 0s, the destination database is being synchronized with the source database in real time. You can also view the data consistency on the **Migration Comparison** tab.

    Figure 2-35 Viewing the synchronization delay

    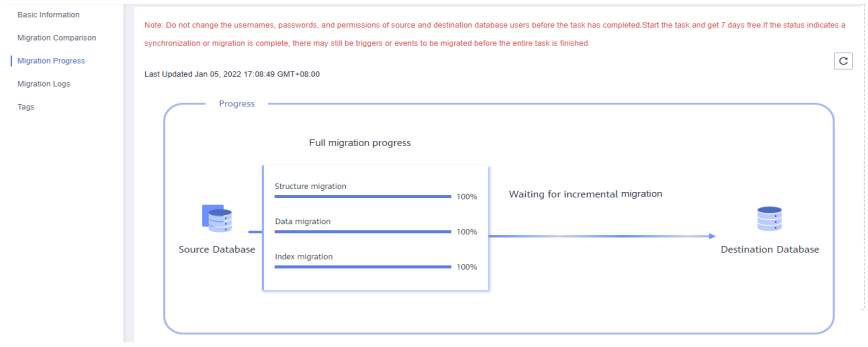

  - Viewing the migration results: On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Comparison** and perform a migration comparison in accordance with the comparison process, which should help you determine an appropriate time for migration to minimize service downtime.

    Figure 2-36 Database comparison process

    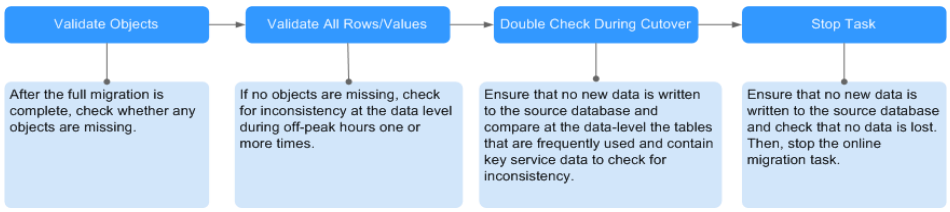

    For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. 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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.

 Copyright © Huawei Technologies Co., Ltd.

2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.

```
show processlist
```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS replication instance. If no additional session executes SQL statements, the service has been stopped.

3. On the **Migration Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the migration is complete.

Step 4 Stop or delete the migration task.

1. Stopping the migration task. After databases and services are migrated to the destination database, to prevent operations on the source database from being synchronized to the destination database to overwrite data, you can stop the migration task. This operation only deletes the replication instance, and the migration task is still displayed in the task list. You can view or delete the task. After the task is stopped, DRS will not charge you anymore.
2. Delete the migration task. After the migration task is complete, you can delete it. After the migration task is deleted, it will no longer be displayed in the task list.

----End

2.5 From On-Premises MongoDB to DDS

DRS helps you migrate data from on-premises MongoDB databases to DDS on the current cloud. With DRS, you can migrate databases online with zero downtime and your services and databases can remain operational during migration.

This section describes how to use DRS to migrate an on-premises MongoDB database to DDS on the current cloud. The following network types are supported:

- VPN
- Public network

 Copyright © Huawei Technologies Co., Ltd.

Diagram

Figure 2-37 VPN

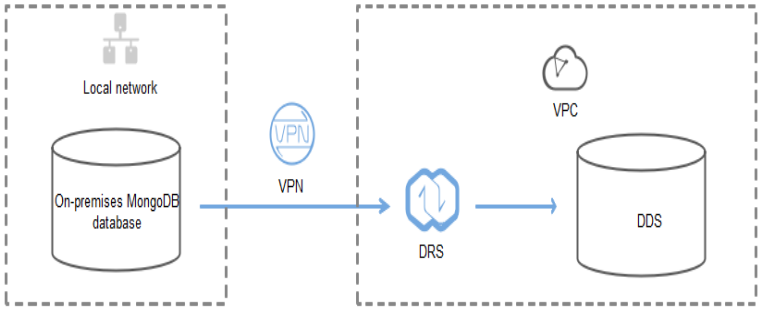

Figure 2-38 Public network+SSL connection

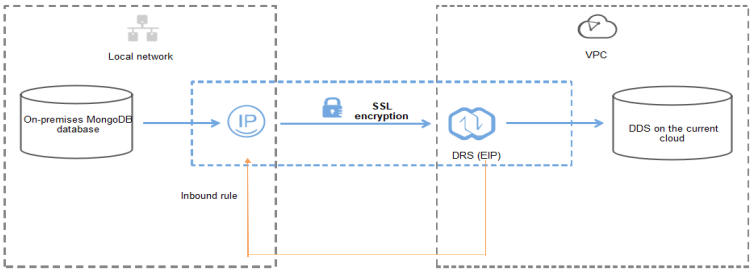

Migration Process

Figure 2-39 Flowchart

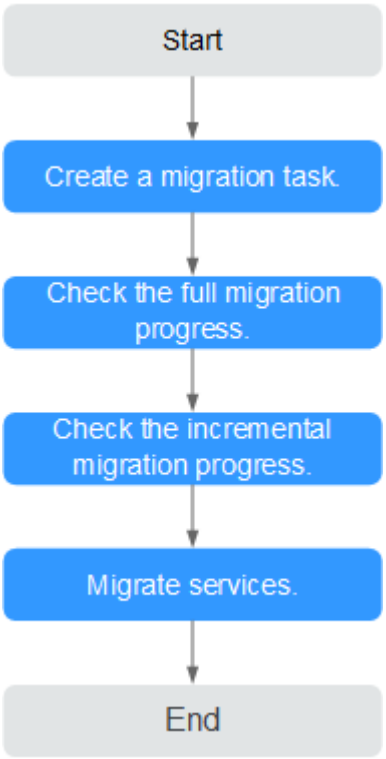

 Copyright © Huawei Technologies Co., Ltd.

Migration Suggestions (Important)

- Database migration is closely impacted by a wide range of environmental and operational factors. To ensure the migration goes smoothly, perform a test run before the actual migration to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.
- It is strongly recommended that you 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.

Notes on Migration (Important)

NOTICE

Before creating a migration task, read the migration notes carefully.

For details, see **precautions** on using specific migration tasks in *Data Replication Service Real-Time Migration*.

Preparations

1. Permissions

   Table 2-24 lists the permissions required for the source and destination databases when migrating data from on-premises MongoDB databases to DDS DB instances.

 Copyright © Huawei Technologies Co., Ltd.

Table 2-24 Migration permissions

Database	Full Migration Permission	Full+Incremental Migration Permission
Source	• Replica set: The source database user must have the read permission for the database to be migrated. • Single node: The source database user must have the read permission for the database to be migrated. • Cluster: The source database user must have the read permission for the databases to be migrated and the config database. • To migrate accounts and roles of the source database, the source database user must have the read permission for the **system.users** and **system.roles** system tables of the admin database.	• Replica set: The source database user must have the read permission for the databases to be migrated and the local database. • Single node: The source database user must have the read permission for the databases to be migrated and the local database. • Cluster: The source mongos node user must have the readAnyDatabase permission for the databases to be migrated and 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 database user must have the read permission for the **system.users** and **system.roles** system tables of the admin database.
Destination	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 migration account must have the read permission for the config database.	

 Copyright © Huawei Technologies Co., Ltd.

- Source database permissions:

  The source database user must have all the required permissions listed in **Table 2-24**. If the permissions are insufficient, create a user that has all of the permissions on the source database.

- Destination database permissions:

  If the destination database is a DDS database, the initial account can be used.

2. Network settings
   - Source database network settings:

     You can migrate on-premises MongoDB databases to DDS through a VPN or public network. Enable public accessibility or establish a VPN for local MongoDB databases based on the site requirements. You are advised to migrate data through a public network, which is more convenient and cost-effective.

   - Destination database network settings:

     - If the source database accesses the destination database through a VPN, enable the VPN service first so that the source database can communicate with the destination DDS network.
     - If you access the DDS DB instance through a public network, no network settings are required.

3. Security rules
   a. Source database network settings:

      - The replication instance needs to be able to access the source DB. That means that the EIP of the replication instance must be on the whitelist of the source MongoDB instance. Before configuring the network whitelist for the source database, you need to obtain the EIP of the DRS replication instance.

        After creating a replication instance on the DRS console, you can find the EIP on the **Configure Source and Destination Databases** page as shown in **Figure 2-40**.

        Figure 2-40 EIP of the replication instance

        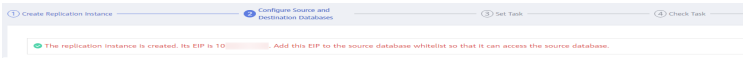

        You can also add 0.0.0.0/0 to the source database whitelist to allow any IP address to access the source database but this action may result in security risks.

      - If the migration is performed over a VPN network, add the private IP address of the DRS replication instance to the whitelist of the source database to enable the source database to communicate with the destination database.

      If you do take this step, then once the migration is complete, you should delete this item from the whitelist or your system will insecure.

 Copyright © Huawei Technologies Co., Ltd.

b. Destination database security group settings:

By default, the destination database and the DRS replication instance are in the same VPC and can communicate with each other. No further configuration is required.

4. Other

You need to export the user information of the MongoDB database first and manually add it to the destination DDS DB instance because the user information will not be migrated.

Migration Procedure

The following describes how to use DRS to migrate an on-premises MongoDB database to a DDS DB instance.

Step 1 Create a migration task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.
2. On the **Online Migration Management** page, click **Create Migration Task**.
3. On the **Create Replication Instance** page, configure the task details, recipient, and replication instance and click **Next**.

Figure 2-41 Replication instance information

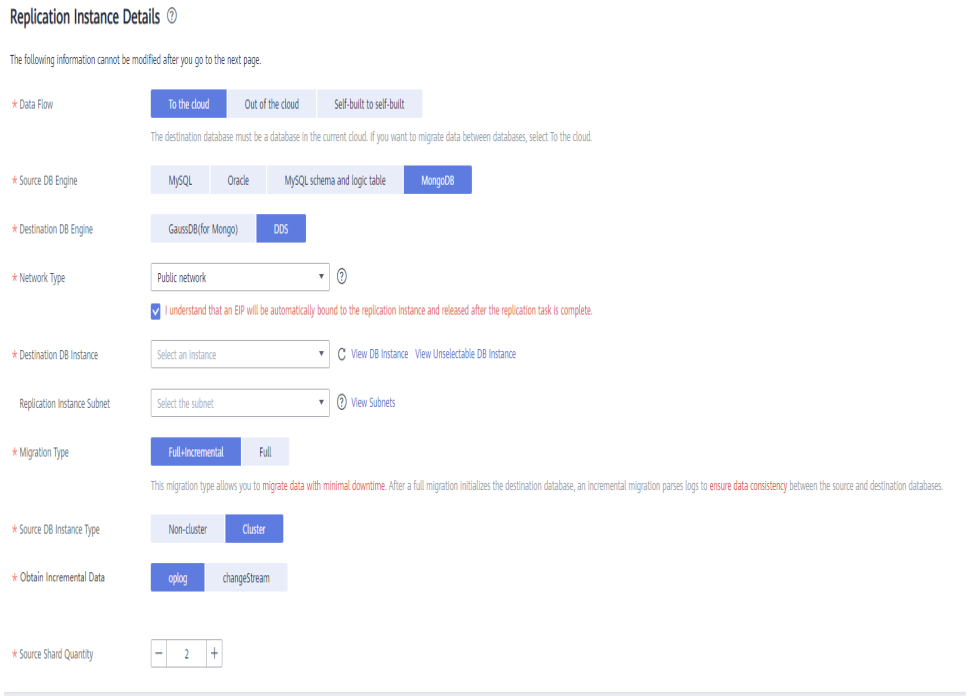

Table 2-25 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.

 Copyright © Huawei Technologies Co., Ltd.

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 consists of a maximum of 256 characters and cannot contain the following special characters: =<>&'\"

Table 2-26 Replication instance settings

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	**MongoDB database**
Destination DB Engine	**DDS**
Network Type	**Public network** Enabling SSL is recommended. It may slow down the migration by 20% to 30% but it ensures data security.
Destination DB Instance	The DDS DB instance you purchased.
Migration Type	– **Full** It migrates all data at one time. If you perform a full migration, you are advised to stop 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** An incremental migration can keep data consistency after a full migration is complete. You can select **Full+Incremental** to enable the source and destination database services to remain accessible during the migration.
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**.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
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 mode, you must disable the balancer of the source instance. When testing the connection, you need to enter the connection information of each shard node of the source instance. – changeStream: This method is recommended. For MongoDB 4.0 and later, DRS connects to mongos nodes of the source instance to extract data. If you select this method, you must enable the WiredTiger storage engine of the source instance. **NOTE** Only whitelisted users can use **changeStream**. To use this function, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets** > **Create Service Ticket** to submit a service ticket.
Source Shard Quantity	If **Source DB Instance Type** is set to **Cluster** and **Obtain Incremental Data** is set to **oplog**, enter the number of source shard nodes. The default minimum number of source DB instances is 2 and the maximum number is 32. You can set this parameter based on the number of source database shards.

4. 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 2-42 Source database information

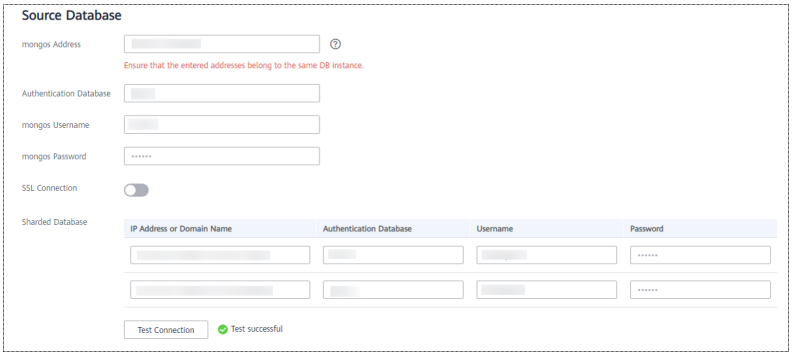

 Copyright © Huawei Technologies Co., Ltd.

Table 2-27 Source database settings

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	A username for the source database.
mongos Password	The password for the source database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.
Sharded Database	Enter the information about the sharded databases in the source database.

- Destination database configuration

  Figure 2-43 Destination database information

 Copyright © Huawei Technologies Co., Ltd.

Table 2-28 Destination database settings

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

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

Figure 2-44 Migration object

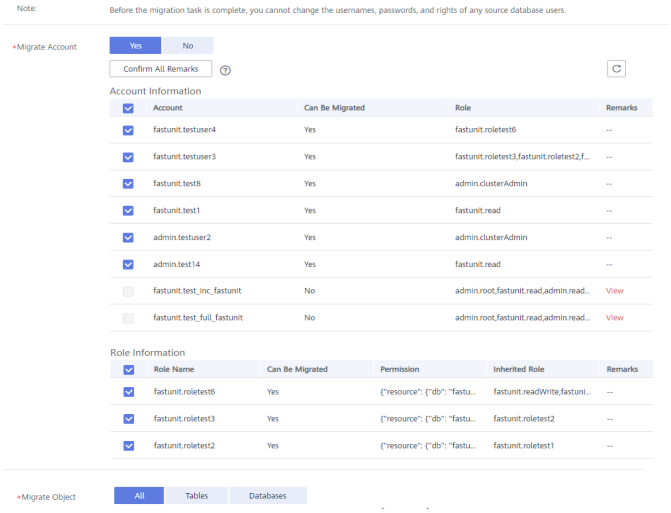

Table 2-29 Migration object

Parameter	Description
Migrate Account	There are accounts that can be migrated completely 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 choose to migrate accounts, see **Migrating Accounts** in *Data Replication Service User Guide* to migrate database users and roles. – **No** During the migration, accounts and roles are not migrated.

 Copyright © Huawei Technologies Co., Ltd.

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 database to be migrated may have dependencies on other objects that are not migrated. To ensure a successful migration, you are advised to 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.

6. 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*.
   - If all check items are successful, click **Next**.

Copyright © Huawei Technologies Co., Ltd.

Figure 2-45 Task Check

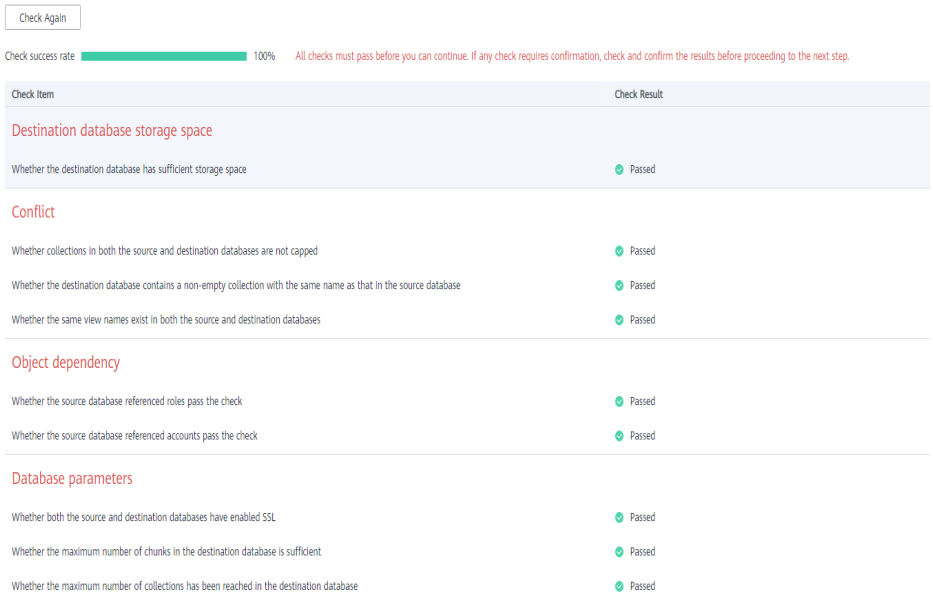

NOTE

You can proceed to the next step only when all check items are successful. If any alarms are generated, view and confirm the alarm details first before proceeding to the 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 2-46 Task startup settings

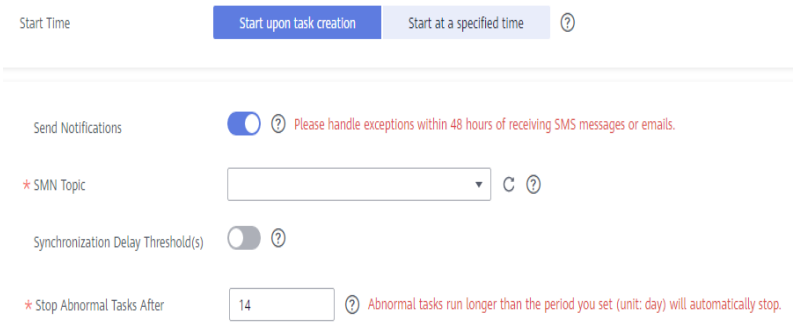

 Copyright © Huawei Technologies Co., Ltd.

Table 2-30 Task startup settings

Parameter	Description
Start Time	Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements. 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 Notification** 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.

8. After the task is submitted, go back to the **Online Migration Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full migration

 Copyright © Huawei Technologies Co., Ltd.

  - Viewing the migration progress: Click the target full migration task, and on the **Migration Progress** tab, you can see the migration progress of the structure, data, indexes, and migration objects. When the progress reaches 100%, the migration is complete.
  - Viewing migration details: 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. Currently, this function is available only to whitelisted users. You can submit a service ticket to apply for this function.
- Incremental Migration Permission
  - Viewing the synchronization delay: After the full migration is complete, an incremental migration starts. On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Progress** to view the synchronization delay of the incremental migration. If the synchronization delay is 0s, the destination database is being synchronized with the source database in real time. You can also view the data consistency on the **Migration Comparison** tab.

    Figure 2-47 Viewing the synchronization delay

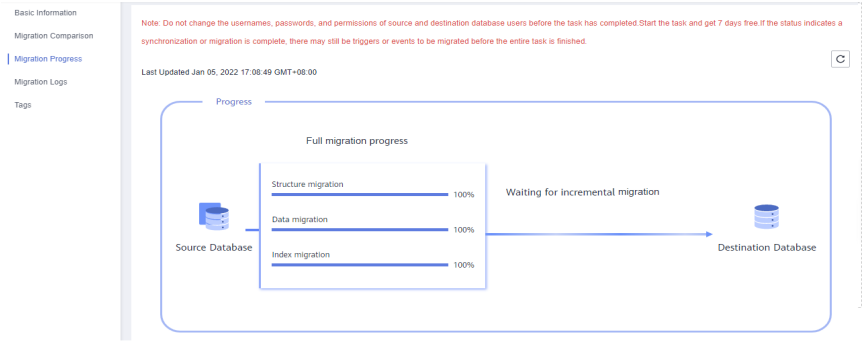

  - Viewing the migration results: On the **Online Migration Management** page, click the target migration task. On the displayed page, click **Migration Comparison** and perform a migration comparison in accordance with the comparison process, which should help you determine an appropriate time for migration to minimize service downtime.

    Figure 2-48 Database comparison process

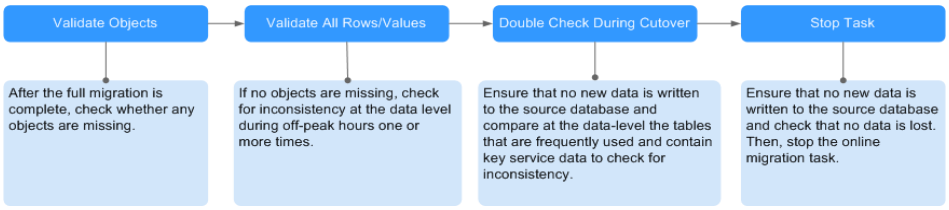

    For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. To obtain accurate

 Copyright © Huawei Technologies Co., Ltd.

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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.
2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.

```
db.currentOp()
```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS replication instance. If no additional session executes SQL statements, the service has been stopped.

3. On the **Migration Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the migration is complete.

Step 4 Stop or delete the migration task.

1. Stopping the migration task. After databases and services are migrated to the destination database, to prevent operations on the source database from being synchronized to the destination database to overwrite data, you can stop the migration task. This operation only deletes the replication instance, and the migration task is still displayed in the task list. You can view or delete the task. DRS will not charge for this task after you stop it.
2. Delete the migration task. After the migration task is complete, you can delete it. After the migration task is deleted, it will no longer be displayed in the task list.

----End

 Copyright © Huawei Technologies Co., Ltd.

3 Backup Migration

3.1 Migrating Microsoft SQL Server Backup Data to RDS SQL Server DB Instance

3.1.1 Overview

It often becomes necessary to hide a real IP address of your database for the sake of security. Migrating data through direct connections is an option, but costly. DRS supports backup migration, which allows you to export data from your source database for backup and upload the backup files to OBS. Then, you can restore the backup files to the destination database to complete the migration. Using this method, data migration can be realized without exposing your source databases.

DRS supports full migration and full+incremental migration.

Scenario 1: Full Backup Migration

In this scenario, you need to stop services, upload the exported full backup file of the Microsoft SQL Server database to OBS, and then restore the backup data to the destination database.

Figure 3-1 Full migration

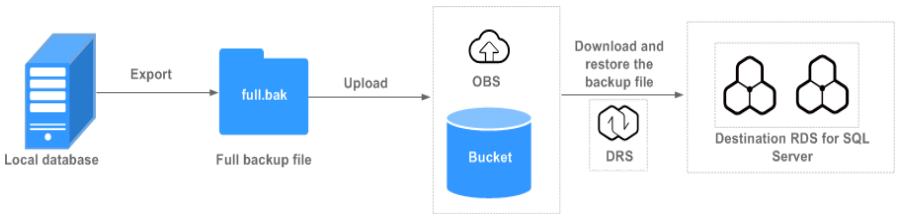

Scenario 2: Full+Incremental Backup Migration

In this scenario, data is migrated continuously. After a full backup and restoration is complete, you can perform incremental migrations for several times to minimize service interruption. A complete restoration process involves restoring several

 Copyright © Huawei Technologies Co., Ltd.

incremental backup files. The destination DB instance will not be available until the last backup file is restored. You can determine whether the selected backup file is the last one.

Figure 3-2 Full+incremental migration

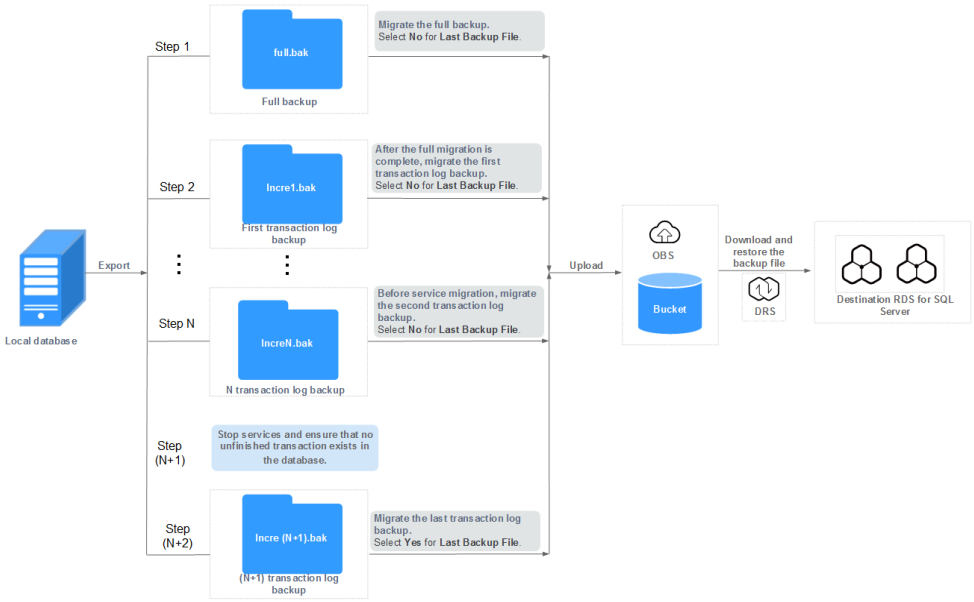

3.1.2 Migration Preparations

This section describes how to prepare for backup and migration. Before using DRS, ensure that you have completed the required preparations.

Preparing for Backup Files

For details, see **Migration Overview** in *Backup Migration*.

Configuring Database Restoration Mode

- No configuration is required during a full backup restoration.
- During full and incremental backup migration, the recovery mode must be set to **Full**. The procedure is as follows:

  Method 1: Log in to the local database center using Microsoft SQL Server Management Studio. Right-click the database to be migrated and choose **Properties** from the shortcut menu. In the dialog box that is displayed, click **Options**, and select **Full** for **Recovery mode**.

Figure 3-3 Recovery mode

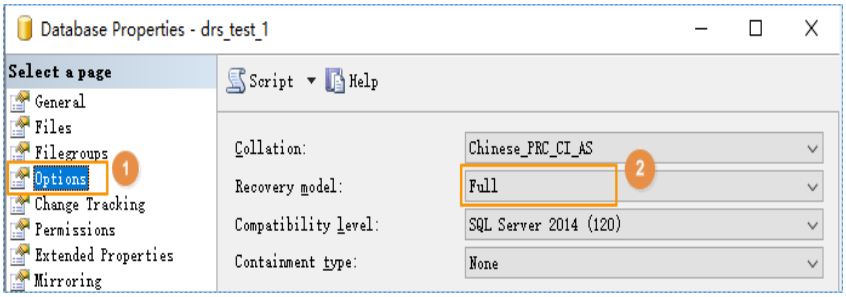

 Copyright © Huawei Technologies Co., Ltd.

Method 2: Run the following SQL commands to set the restoration mode:

```
USE master;
ALTER DATABASE database_name SET RECOVERY FULL;
```

Destination Database Storage Space

The available storage space of the destination database should be at least 1.5 times larger than the total storage space of the database to be restored.

Migration Duration Reference

A complete backup and migration process consists of the following phases:

Figure 3-4 Migration diagram

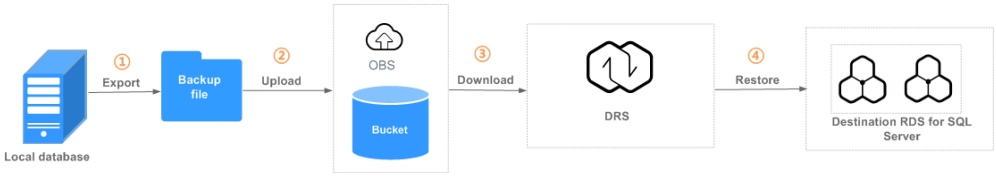

Table 3-1 Backup migration

Phase	Name	Description
①	Export the backup files.	The time required for generating database backup files depends on the configuration of the source database. You need to estimate the time based on the configuration of the source database.
②	Upload the backup file to an OBS bucket.	OBS does not limit the upload and download speed. If you access OBS through a public network, the upload and download speed is restricted by the public network bandwidth. For example, if the bandwidth of the public network is 10 MB/s and the network is not affected by other factors, the upload rate is 10 MB/s.
③	Download the backup file to the destination RDS SQL Server through DRS.	Generally, the download speed is 100 MB/s or 300 GB/h.
④	Restore the backup files to the destination database.	Generally, the recovery speed is about 5 GB/min or 300 GB/h.
Total Duration	Total time consumption = phase ① time consumption + phase ② time consumption + phase ③ time consumption + phase ④ time consumption Service downtime = Service suspension -> Last incremental backup -> Uploading to OBS -> Creating a DRS task for restoration	

 Copyright © Huawei Technologies Co., Ltd.

The following uses an example to describe the time required for backing up and migrating data. You can estimate the migration time in advance. The actual time required depends on the network and database configurations on the client. The time listed in the following table is for reference only.

Example

Table 3-2 Backup migration example

Phase	Name	Backup File Size (GB)	Time Required (h)
①	Export the backup files.	283	5.5
②	Upload the backup file to an OBS bucket.	283	8.95
③	Download the backup file by through DRS.	283	0.61
④	Restore the backup files to the destination database.	283	2.24
Total Duration			17.3

3.1.3 Exporting Backup Files

This section describes how to export the full backup files and incremental backup files of a database.

Step 1 Check the parameter settings of the local database.

NOTE

This operation is mandatory for full and incremental migration. If you only perform the full migration, skip this step.

The database has log truncation and shrinking configurations. Therefore, before exporting full backup files, you must set the database recovery model to **Full** until the entire database and services are migrated to the DB instance on the current cloud.

1. Log in to the local database center through Microsoft SQL Server Management Studio.
2. Right-click the database to be migrated and choose **Properties** from the shortcut menu. In the displayed dialog box, select **Options** from the left list.
3. Select **Full** for **Recovery model**. Then, click **OK**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-5 Configuring the recovery model

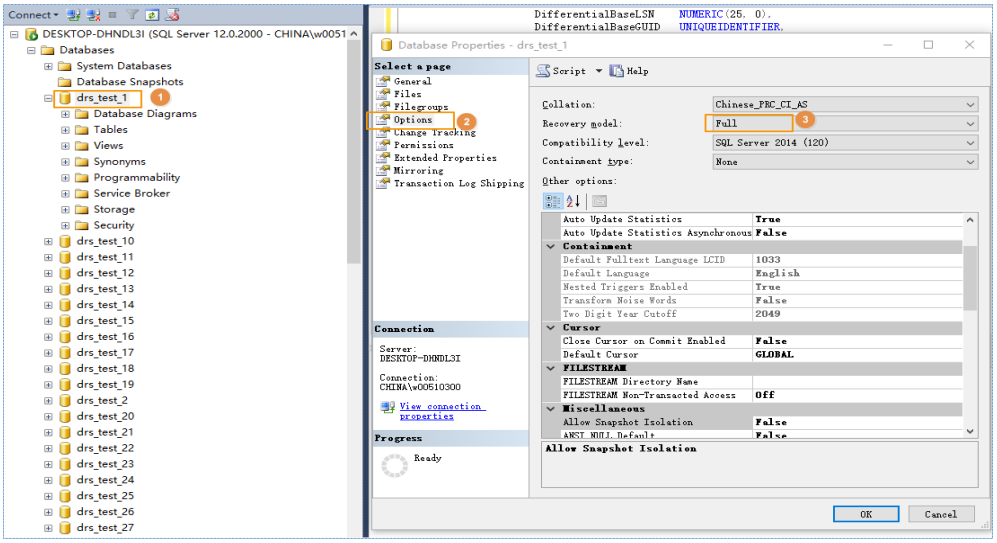

Step 2 (Optional) Set backup file compression parameter.

If the bandwidth of your local data center is low and uploading data to OBS Browser takes a long time, you are advised to configure the backup file compression parameters.

1. Log in to the database center through Microsoft SQL Server Management Studio as the database administrator.
2. In Object Explorer, right-click the server and choose **Properties**.
3. Click **Database Settings**.
4. In the **Backup and restore** area, select **Compress backup**.

   This setting determines the server-level default setting for compression backup, as detailed below:

   - If **Compress backup** is not selected, the new backup will not be compressed by default.
   - If **Compress backup** is selected, the new backup file is compressed by default.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-6 Configuring compression parameter

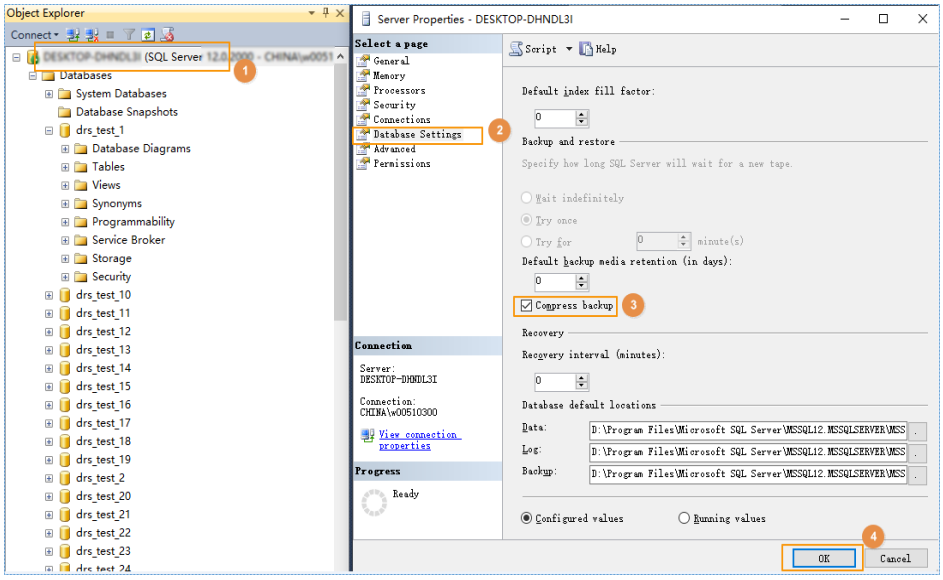

Step 3 Export the full backup file.

1. Log in to the local database center through Microsoft SQL Server Management Studio.
2. Right-click the database whose data needs to be exported and choose **Tasks** > **Back Up**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-7 Back Up

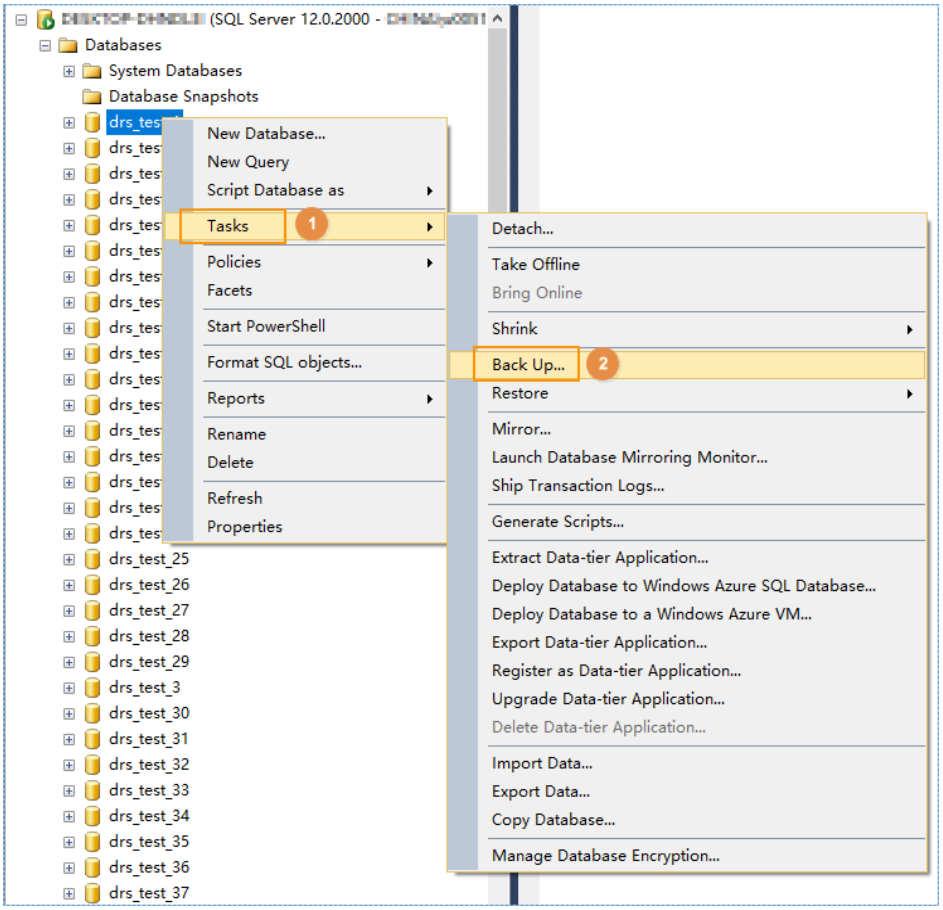

3. Set **Backup type** to **Full**, click **Add**, and enter the path to which the backup file is exported. The file name extension must be **.bak**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-8 Setting the full backup file

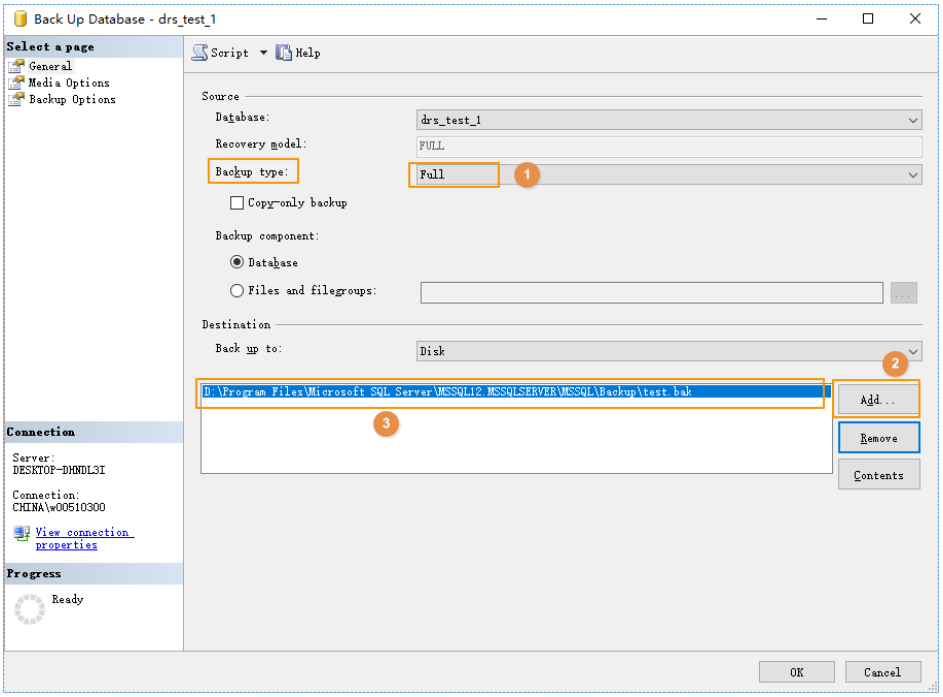

NOTE

- It is recommended that the backup file name be the same as the database name (case sensitive), and the suffix **.bak** be added to the backup file name.
- You are advised to back up all the databases in a .bak file or a few .bak files. This reduces the frequency of uploading and restoring the database.

Step 4 Export the incremental backup file.

1. Log in to the local database center through Microsoft SQL Server Management Studio.
2. Right-click the database whose data needs to be exported and choose **Tasks** > **Back Up**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-9 Back Up

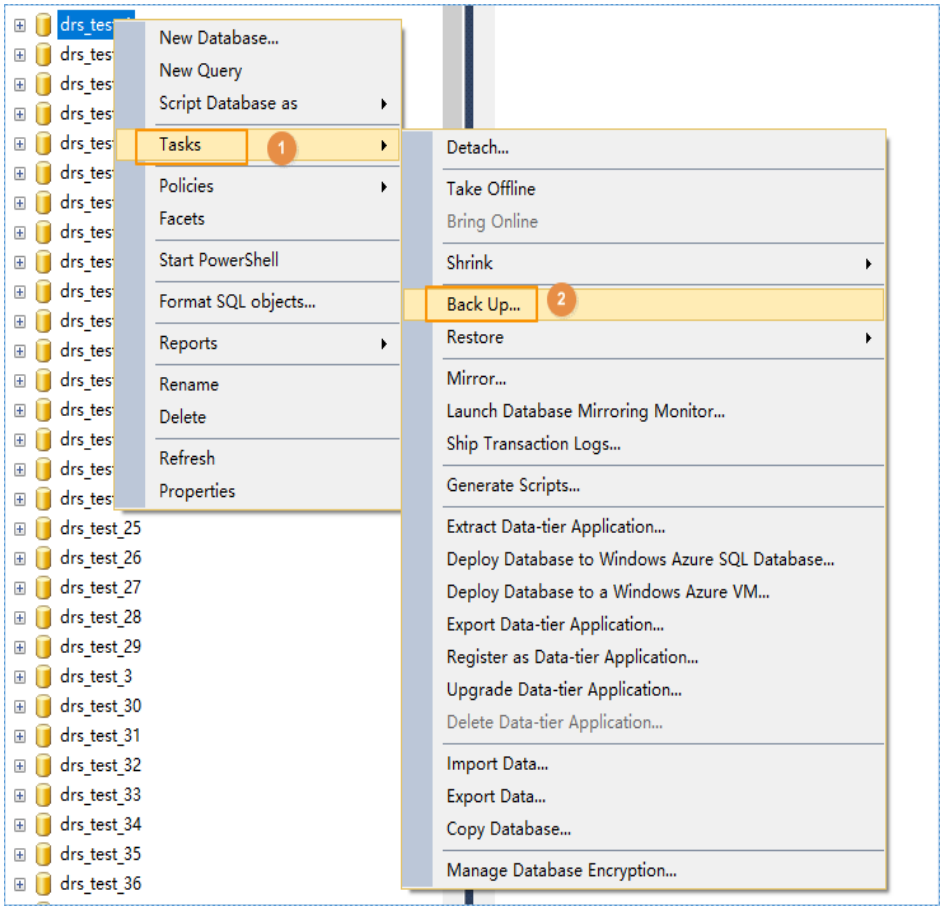

3. Select **Transaction Log** for **Backup type**, click **Add**, and enter the path to which the backup file is exported. The suffix is **.bak**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-10 Setting the incremental backup file

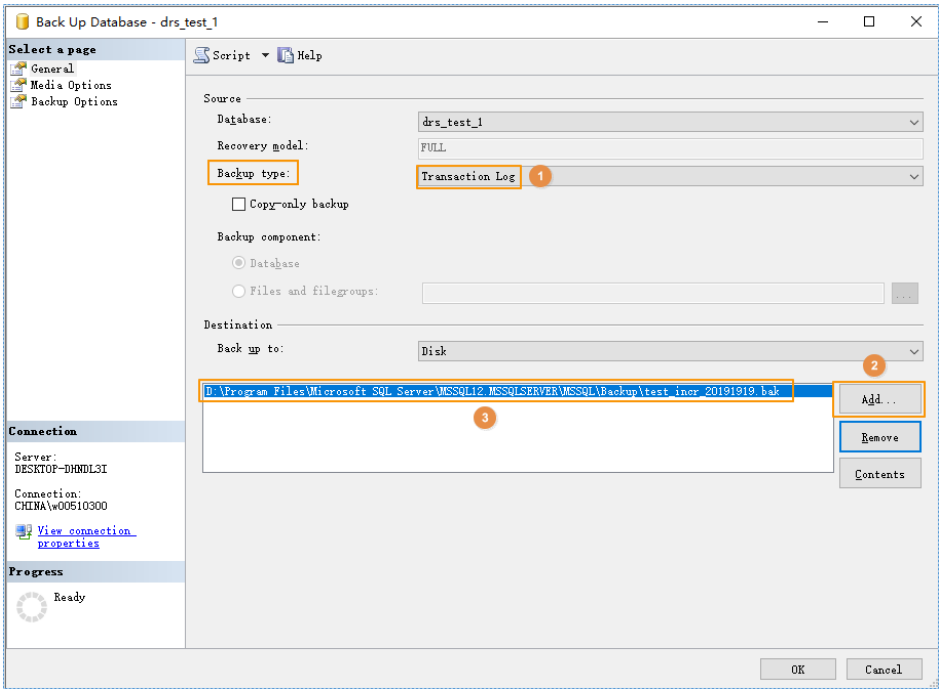

NOTE

- It is recommended that the backup file name be the same as the database name (case sensitive), with the time stamp and .bak suffix. For example: [*Database name*] _**Incr**_ [*Timestamp*] **.bak**.
- You are advised to back up all the databases in a .bak file or a few .bak files. This reduces the frequency of uploading and restoring the database. For example, you can back up the A, B, and C databases to a .bak file, and then compress, upload, and restore the three databases in a batch. This helps to improve the success rate of data restoration.

----End

3.1.4 Uploading Backup Files

This section describes how to upload backup files.

Step 1 Create an OBS bucket and upload the backup file to the OBS bucket.

 Copyright © Huawei Technologies Co., Ltd.

NOTE

- If the size of the file uploaded at a time does not exceed 5 GB, select **Standard** for **Storage Class** and **Public Read** for **Bucket Policy** on the OBS console.
- If you want to upload multiple files in batches (up to 100 files can be uploaded at a time and the total size of the files cannot exceed 5 GB) or upload files greater than 5 GB at a time, . For details, see .

  You are advised to store backup files in independent OBS buckets in the same region to distinguish the backup files from other public files.

  Before uploading the backup file, create access keys (AK and SK) in the OBS console. For details, see .

  When uploading backup files, select **Standard** for **Storage Class**.
- Microsoft SQL Server supports only .bak files and does not support restoration from a newer version to an older version.
- You are advised to store backup files in independent OBS buckets in the same region to distinguish the backup files from other public files.

1. Log in to OBS Console..
2. For details about how to add an OBS bucket, see **Creating a Bucket**.
3. For details about how to upload files to an OBS bucket, see **Uploading a File**.

----End

3.1.5 Scenario 1: Full Backup Migration

In this scenario, you need to stop services, upload the exported full backup file of the Microsoft SQL Server database to OBS, and then restore the backup data to the destination database.

This section describes how to migrate a full backup.

Migration Diagram

Figure 3-11 Full migration

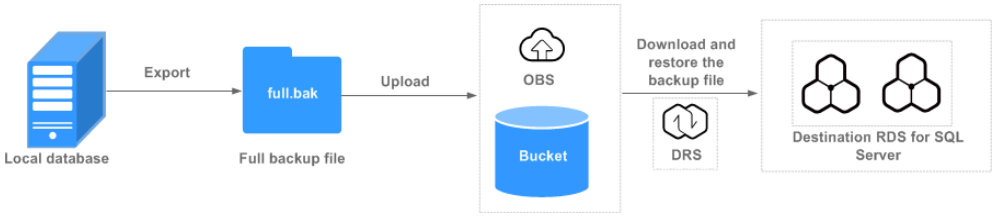

 Copyright © Huawei Technologies Co., Ltd.

Migration Process

Figure 3-12 Migration process

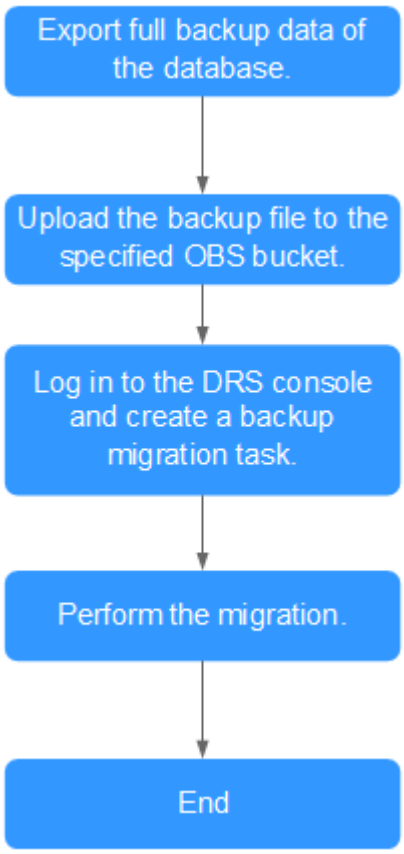

Procedure

Step 1 Export the full backup file of the database. For details, see **Exporting Backup Files**.

Step 2 Upload files to an OBS bucket. For details, see **Uploading Backup Files**.

Step 3 Log in to the DRS console.

Step 4 Click in the upper left corner and select a region and a project.

Step 5 Choose **Service List** > **Databases** > **Data Replication Service**.

Step 6 In the navigation pane on the left, choose **Backup Migration Management**. Then, click **Create Migration Task**.

Step 7 Enter the information about the migration task and backup file, and click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-13 Migration task information

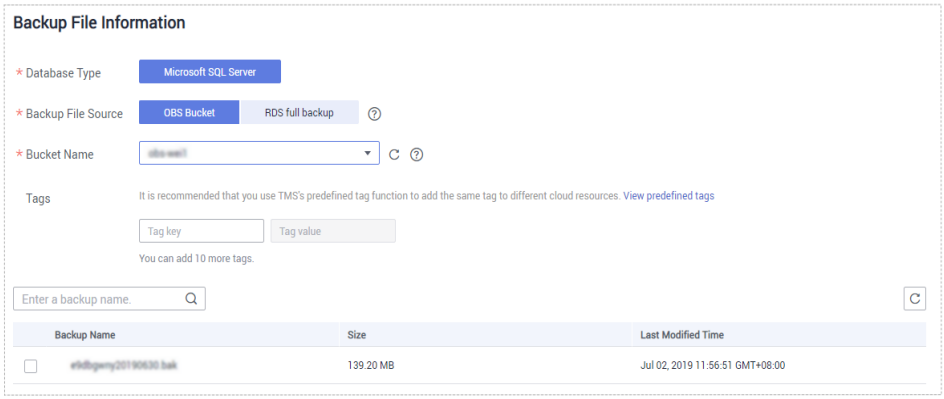

Table 3-3 Migration task information

Parameter	Description
Task Name	The task name consists of 4 to 64 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: !=<>'&"
Database Type	Select **Microsoft SQL Server**.
Backup File Source	Select **OBS bucket**.
Bucket Name	Select the bucket where the backup file is stored and the full backup file. **NOTE** • The backup file of the Microsoft SQL Server must be in the .bak format. Multiple backup files can be selected at the same time. • A database cannot be split into different files for upload. • The bucket name, backup file name, or path cannot contain Chinese characters. • If the number of data records in a bucket exceeds 500, data overflow may occur. As a result, some data cannot be displayed. You can use an independent standard bucket with the public-read permission.

Step 8 On the **Select Destination** page, specify database information and click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-14 Full migration database information

Table 3-4 Microsoft SQL Server database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. If no RDS DB instance is available, you can buy one by referring to the
Backup File Format	Select **Full**. **Full**: indicates full backup files. **NOTE** To migrate databases at a time, you need to stop services first and upload full backups for restoration.
Last Backup	Select **Yes** if you perform a one-time full migration.
Overwrite Data	You can determine whether to overwrite data in the destination database during the restoration if the destination DB instance contains a database with the same name as the backup database. **NOTE** If you select this option, the destination databases with the same names as the backup databases will be overwritten. Exercise caution when performing this operation.
Perform Pre-verification	Specifies whether to perform pre-verification on the backup migration task. The default value is **Yes**. • **Yes**: To ensure successful migration and identify potential problems in advance, verify the validity, integrity, continuity, and version compatibility of backup files before restoration. • **No**: If pre-verification is not performed, the migration speed is faster, but you need to check the validity, integrity, continuity, and version compatibility of backup files on your own.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Restore Database	You can restore all or some of databases. • **All**: Restores all databases in the backup file. You do not need to enter the names of the databases to be restored. By default, all databases in the backup file are restored. • **Custom**: Restore specified databases from the backup file. You need to enter the names of the databases to be restored. The databases to be restored must be consistent with those for which the full or incremental backups are created.
Reset Database Name	If **Restore Database** is set to **All**, you can reset database names. If you enable **Reset Database Name**, the original database names in the backup file will be reset to new database names. Prerequisites: • The backup file contains only one database. • The backup file is a full backup file and is the last backup file. Select **Full** for **Backup File** and **Yes** for **Last Backup Type**. **NOTE** The database name can be reset only when **Backup Type** is set to **Full** and **Restore Database** is set to **All**.
Backup Database Name	If **Restore Database** is set to **Custom**, specify **Backup Database Name**. The backup database name is case sensitive and must be the same as that in the backup file. The backup database name can contain a maximum of 256 bytes, and the new database name can contain a maximum of 128 bytes. Only letters, digits, hyphens (-), and underscores (_) are allowed. DRS also allows you to set an alias for the database to be restored.

Step 9 On the **Confirm Task** page, check the configuration details, read and agree to the agreement, and click **Next**.

Step 10 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

3.1.6 Scenario 2: Full+Incremental Backup Migration

In this scenario, data is migrated continuously. After a full backup and restoration is complete, you can perform incremental migrations for several times to minimize service interruption. A complete restoration process involves restoring several incremental backup files. The destination DB instance will not be available until

 Copyright © Huawei Technologies Co., Ltd.

the last backup file is restored. You can determine whether the selected backup file is the last one.

This section uses a full+incremental migration as an example to describe how to perform a migration while minimizing service interruption.

Migration Diagram

Figure 3-15 Full+incremental migration

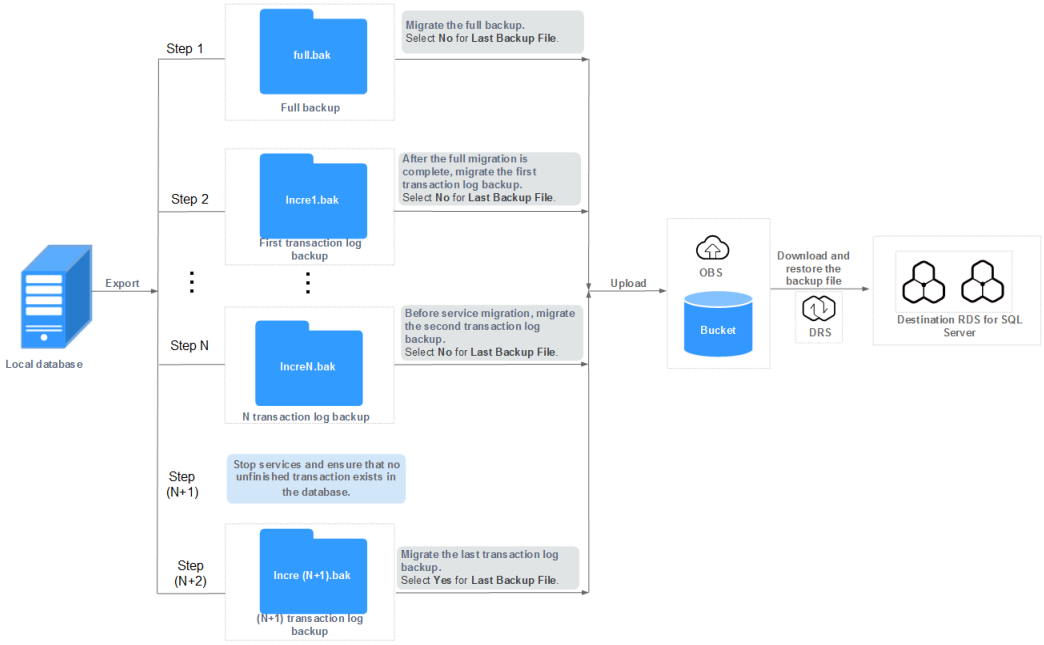

Migration Process

Figure 3-16 Flowchart

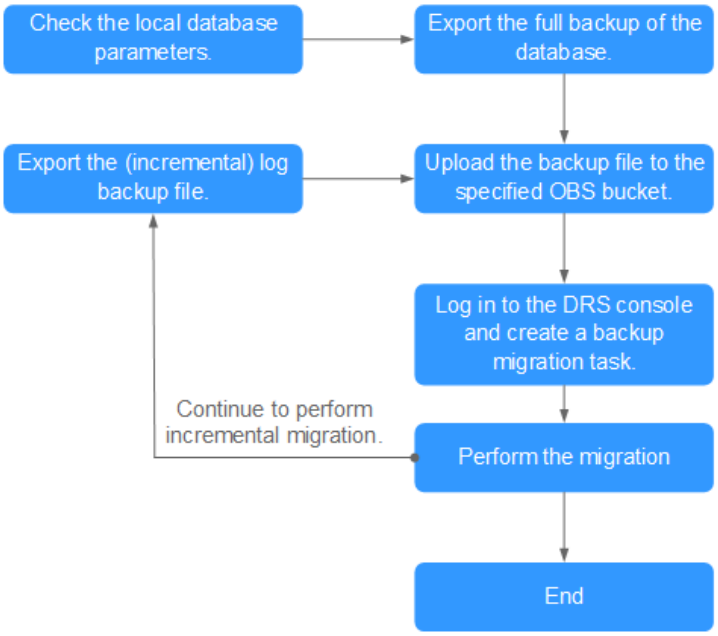

 Copyright © Huawei Technologies Co., Ltd.

First Full Migration

Step 1 Export the full backup file of the database. For details, see **Exporting Backup Files**.

Step 2 Upload files to an OBS bucket. For details, see **Uploading Backup Files**.

Step 3 Log in to the DRS console.

Step 4 Click in the upper left corner and select a region and a project.

Step 5 Choose **Service List** > **Databases** > **Data Replication Service**.

Step 6 In the navigation pane on the left, choose **Backup Migration Management**. Then, click **Create Migration Task**.

Step 7 Enter the information about the migration task and backup file, and click **Next**.

Figure 3-17 Migration task information

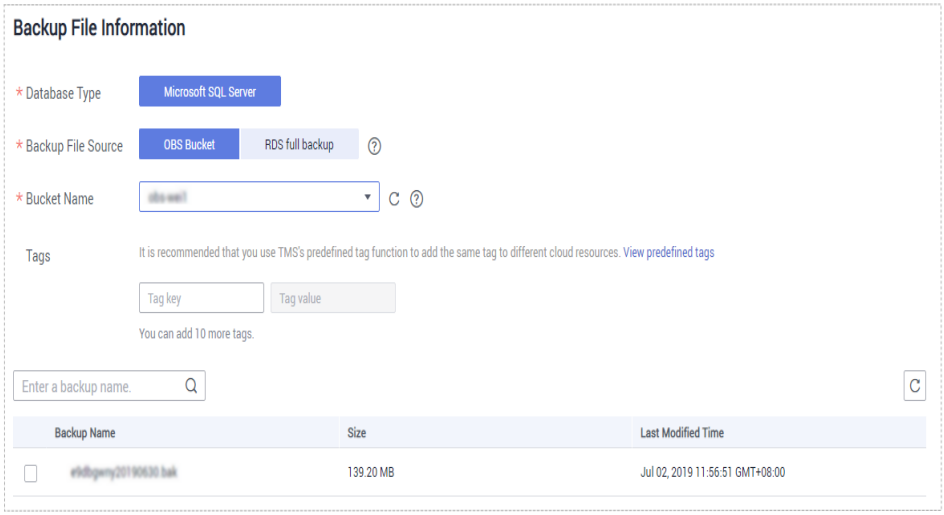

Table 3-5 Migration task information

Parameter	Description
Task Name	The task name consists of 4 to 64 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: !=<>'&"
Database Type	Select **Microsoft SQL Server**.
Backup File Source	Select **OBS bucket**.

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Bucket Name	Select the bucket where the backup file is stored and the full backup file. **NOTE** • The backup file of the Microsoft SQL Server must be in the .bak format. Multiple backup files can be selected at the same time. • A database cannot be split into different files for upload. • The bucket name, backup file name, or path cannot contain Chinese characters. • If the number of data records in a bucket exceeds 500, data overflow may occur. As a result, some data cannot be displayed. You can use an independent standard bucket with the public-read permission.

Step 8 On the **Select Destination** page, specify database information and click **Next**.

Figure 3-18 Full migration database information

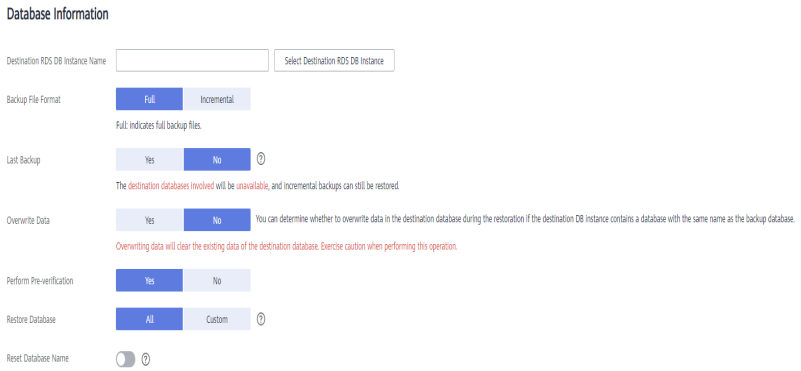

Table 3-6 Microsoft SQL Server database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. If no RDS DB instance is available, you can buy one by referring to the
Backup Type	Select **Full**. **Full**: indicates full backup files.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Last Backup File	A complete restoration process involves restoring several incremental backup files. The destination DB instance will not be available until the last backup file is restored. You can determine whether the selected backup file is the last one. Select **Yes** in either of the following scenarios: ● Perform a one-time full migration. ● The selected backup file is the last one to be restored. Select **No** in the scenario where you continue to restore databases using incremental backup files after a full backup restoration is performed. In this case, the destination database is in the restoring state and cannot be read or written.
Overwrite Data	You can determine whether to overwrite data in the destination database during the restoration if the destination DB instance contains a database with the same name as the backup database. **NOTE** If you select this option, the destination databases with the same names as the backup databases will be overwritten. Exercise caution when performing this operation.
Perform Pre-verification	Specifies whether to perform pre-verification on the backup migration task. The default value is **Yes**. ● **Yes**: To ensure successful migration and identify potential problems in advance, verify the validity, integrity, continuity, and version compatibility of backup files before restoration. ● **No**: If pre-verification is not performed, the migration speed is faster, but you need to check the validity, integrity, continuity, and version compatibility of backup files on your own.
Restore Database	You can restore all or some of databases. ● **All**: Restores all databases in the backup file. You do not need to enter the names of the databases to be restored. By default, all databases in the backup file are restored. ● **Custom**: Restore specified databases from the backup file. You need to enter the names of the databases to be restored. The databases to be restored must be consistent with those for which the full or incremental backups are created.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Reset Database Name	If **Restore Database** is set to **All**, you can reset database names. If you enable **Reset Database Name**, the original database names in the backup file will be reset to new database names. Prerequisites: • The backup file contains only one database. • The backup file is a full backup file and is the last backup file. Select **Full** for **Backup File** and **Yes** for **Last Backup Type**. **NOTE** The database name can be reset only when **Backup Type** is set to **Full** and **Restore Database** is set to **All**.
Backup Database Name	If **Restore Database** is set to **Custom**, specify **Backup Database Name**. The backup database name is case sensitive and must be the same as that in the backup file. The backup database name can contain a maximum of 256 bytes, and the new database name can contain a maximum of 128 bytes. Only letters, digits, hyphens (-), and underscores (_) are allowed. DRS also allows you to set an alias for the database to be restored.

Step 9 On the **Confirm Task** page, check the configuration details, read and agree to the agreement, and click **Next**.

Step 10 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

First Incremental Migration

Step 1 Export the first incremental backup file of the database. For details, see **Exporting Backup Files**.

Step 2 Upload files to an OBS bucket. For details, see **Uploading Backup Files**.

Step 3 Log in to the DRS console.

Step 4 In the navigation pane on the left, choose **Backup Migration Management**. Then, click **Create Migration Task**.

Step 5 Enter the information about the migration task and backup file, and click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-19 Incremental backup

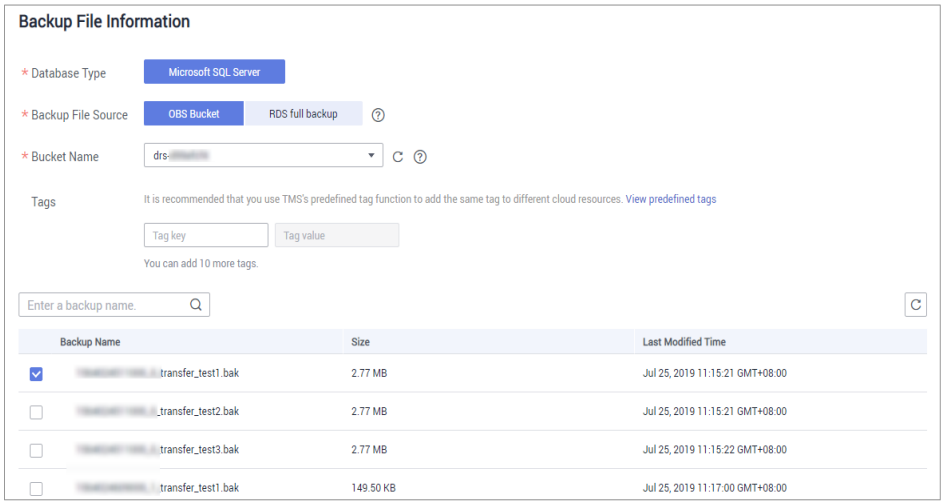

Table 3-7 Migration task information

Parameter	Description
Task Name	The task name consists of 4 to 64 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: !=<>'&"
Database Type	Select **Microsoft SQL Server**.
Backup File Source	Select **OBS bucket**.
Bucket Name	Select the bucket where the backup file is stored and the first incremental backup file. **NOTE** • The backup file of the Microsoft SQL Server must be in the .bak format. Multiple backup files can be selected at the same time. • A database cannot be split into different files for upload. • The bucket name, backup file name, or path cannot contain Chinese characters.

Step 6 On the **Select Destination** page, specify database information and click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-20 Incremental migration database information

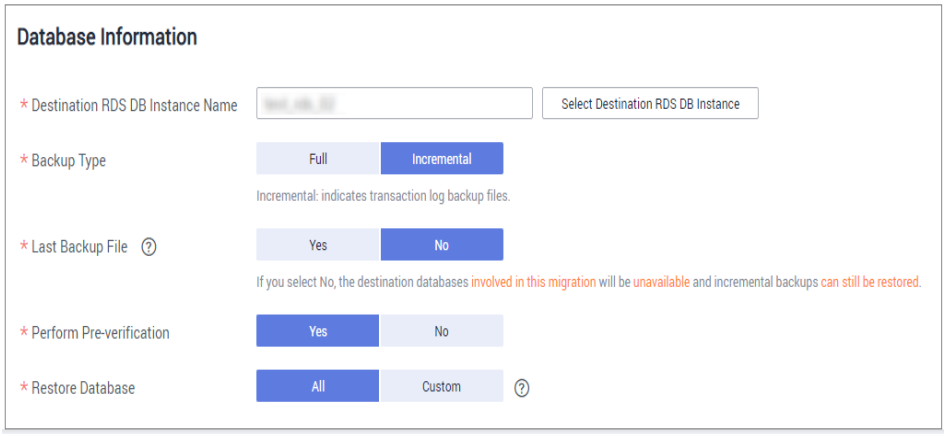

Table 3-8 Microsoft SQL Server database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. The destination RDS DB instance must be the same as the DB instance selected during full backup and restoration.
Backup Type	Select **Incremental**. **Incremental**: indicates log backup files.
Last Backup File	A complete restoration process involves restoring several incremental backup files. The destination DB instance will not be available until the last backup file is restored. You can determine whether the selected backup file is the last one. Select **Yes** in either of the following scenarios: • Perform a one-time full migration. • The selected backup file is the last one to be restored. Select **No** if the incremental backup is performed for the first time. In this case, the destination database is in the restoring state and cannot be read or written.
Perform Pre-verification	Specifies whether to perform pre-verification on the backup migration task. The default value is **Yes**. • **Yes**: To ensure successful migration and identify potential problems in advance, verify the validity, integrity, continuity, and version compatibility of backup files before restoration. • **No**: If pre-verification is not performed, the migration speed is faster, but you need to check the validity, integrity, continuity, and version compatibility of backup files on your own.

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Restore Database	You can restore all or some of databases. • **All**: Restores all databases in the backup file. You do not need to enter the names of the databases to be restored. By default, all databases in the backup file are restored. • **Custom**: Restore specified databases from the backup file. You need to enter the names of the databases to be restored.

Step 7 On the **Confirm Task** page, check the configuration details, read and agree to the agreement, and click **Next**.

Step 8 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

Second Incremental Migration

To minimize the service interruption caused by migration, you need to back up, upload, and restore transaction logs before service migration. These operations will restore historical data to the destination database, which greatly reduces the time for uploading and restoring the last transaction log backup during the migration.

Step 1 Before service migration, export new incremental backup files. For details, see **Exporting Backup Files**.

Step 2 Perform **Step 2** to **Step 4**.

Step 3 Enter the information about the migration task and backup file, and click **Next**.

Figure 3-21 Incremental backup

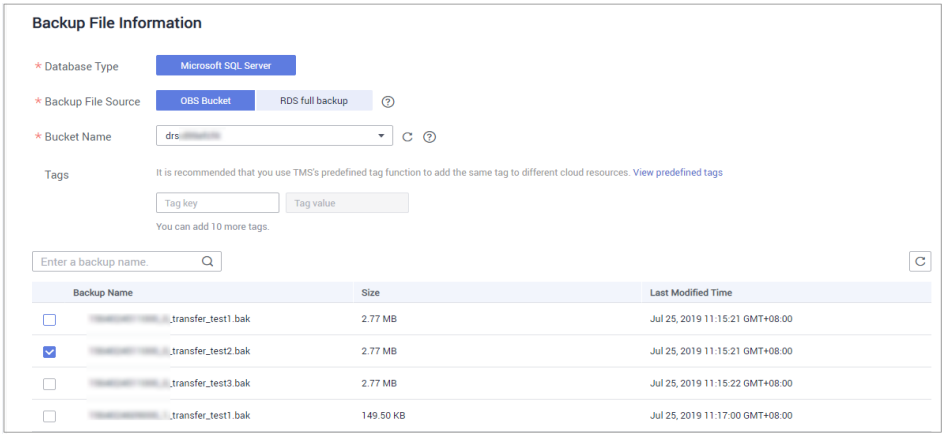

 Copyright © Huawei Technologies Co., Ltd.

Table 3-9 Migration task information

Parameter	Description
Task Name	The task name consists of 4 to 64 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: !<>&'\"
Database Type	Select **Microsoft SQL Server**.
Backup File Source	Select **OBS bucket**.
Bucket Name	Select the bucket where the backup file is stored and the new incremental backup file. **NOTE** • The backup file of the Microsoft SQL Server must be in the .bak format. Multiple backup files can be selected at the same time. • A database cannot be split into different files for upload. • The bucket name, backup file name, or path cannot contain Chinese characters.

Step 4 On the **Select Destination** page, specify database information and click **Next**.

Figure 3-22 Incremental migration database information

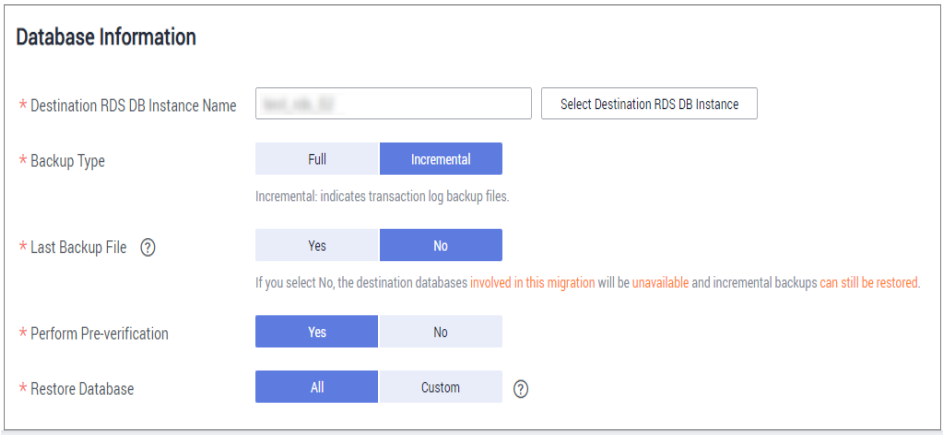

Table 3-10 Microsoft SQL Server database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. The destination RDS DB instance must be the same as the DB instance selected during full backup and restoration.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Backup Type	Select **Incremental**. **Incremental**: indicates log backup files.
Last Backup File	A complete restoration process involves restoring several incremental backup files. The destination DB instance will not be available until the last backup file is restored. You can determine whether the selected backup file is the last one. Select **Yes** in either of the following scenarios: • Perform a one-time full migration. • The selected backup file is the last one to be restored. Select **No** if the incremental backup is performed before migration. In this case, the destination database is in the restoring state and cannot be read or written.
Perform Pre-verification	Specifies whether to perform pre-verification on the backup migration task. The default value is **Yes**. • **Yes**: To ensure successful migration and identify potential problems in advance, verify the validity, integrity, continuity, and version compatibility of backup files before restoration. • **No**: If pre-verification is not performed, the migration speed is faster, but you need to check the validity, integrity, continuity, and version compatibility of backup files on your own.
Restore Database	You can restore all or some of databases. • **All**: Restores all databases in the backup file. You do not need to enter the names of the databases to be restored. By default, all databases in the backup file are restored. • **Custom**: Restore specified databases from the backup file. You need to enter the names of the databases to be restored.

Step 5 On the **Confirm Task** page, check the configuration details, read and agree to the agreement, and click **Next**.

Step 6 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

Checking Database Transactions

Before service migration, stop services and ensure that no unfinished transaction exists in the database. Otherwise, data may be lost during migration.

Step 1 Run the following statement to check whether the IP addresses of the service system are disconnected:

 Copyright © Huawei Technologies Co., Ltd.

```
select * from sys.dm_exec_connections;
```

- If yes, all service system IP addresses are disconnected and the last incremental backup can be migrated.
- If no, go to **Step 2**.

Step 2 If an IP address is not disconnected, run the following statement to query the sessions that are not closed:

```
select * from sys.dm_exec_sessions;
```

Run the following statement to view the transactions that are being executed:

```
select * from sys.dm_tran_session_transactions;
```

If the query result contains unclosed sessions and transactions that are being executed, go to **Step 3**.

Step 3 The last incremental backup can be migrated only after the transaction is complete, the session is closed, and the service system is disconnected.

----End

Last Incremental Migration

After several incremental migrations, the two databases are almost consistent. During the transaction checks, services are stopped and no data will be generated. Then, you need to perform the last incremental migration to ensure the data consistency.

Step 1 Export the new incremental backup file of the database. For details, see **Exporting Backup Files**.

Step 2 Perform **Step 2** to **Step 4**.

Step 3 Enter the information about the migration task and backup file, and click **Next**.

Figure 3-23 Incremental backup

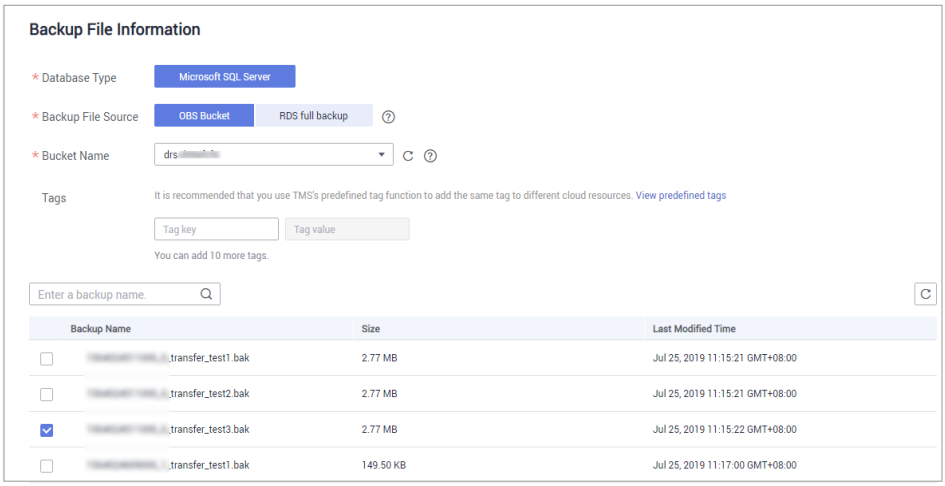

 Copyright © Huawei Technologies Co., Ltd.

Table 3-11 Migration task information

Parameter	Description
Task Name	The task name consists of 4 to 64 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: !=<>'&"
Database Type	Select **Microsoft SQL Server**.
Backup File Source	Select **OBS bucket**.
Bucket Name	Select the bucket where the backup file is stored and the new incremental backup file. **NOTE** • The backup file of the Microsoft SQL Server must be in the .bak format. Multiple backup files can be selected at the same time. • A database cannot be split into different files for upload. • The bucket name, backup file name, or path cannot contain Chinese characters.

Step 4 On the **Select Destination** page, specify database information and click **Next**.

Figure 3-24 Incremental migration database information

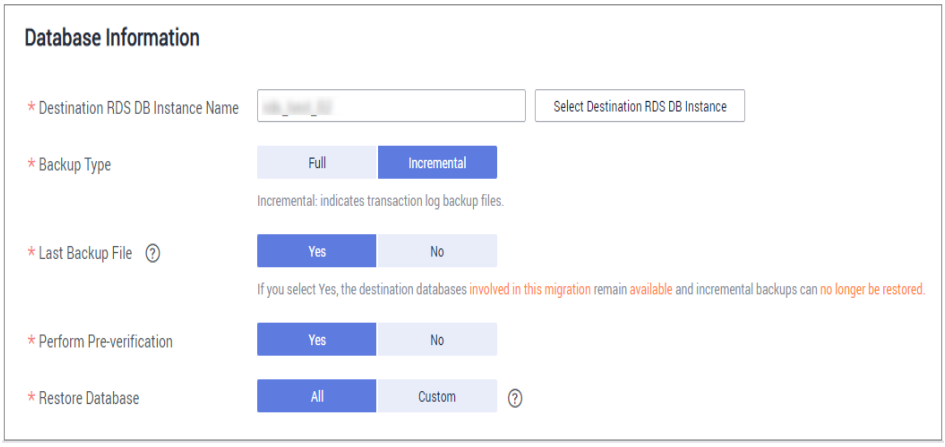

Table 3-12 Microsoft SQL Server database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. The destination RDS DB instance must be the same as the DB instance selected during full backup and restoration.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Backup Type	Select **Incremental**. **Incremental**: indicates log backup files.
Last Backup File	Select **Yes** because it is the last incremental migration after services are stopped.
Perform Pre-verification	Specifies whether to perform pre-verification on the backup migration task. The default value is **Yes**. • **Yes**: To ensure successful migration and identify potential problems in advance, verify the validity, integrity, continuity, and version compatibility of backup files before restoration. • **No**: If pre-verification is not performed, the migration speed is faster, but you need to check the validity, integrity, continuity, and version compatibility of backup files on your own.
Restore Database	You can restore all or some of databases. • **All**: Restores all databases in the backup file. You do not need to enter the names of the databases to be restored. By default, all databases in the backup file are restored. • **Custom**: Restore specified databases from the backup file. You need to enter the names of the databases to be restored.

Step 5 On the **Confirm Task** page, check the configuration details, read and agree to the agreement, and click **Next**.

Step 6 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

3.1.7 Manual Configuration

Scenarios

After data is migrated from the local host or VMs to the RDS SQL Server DB instance on the current cloud through DRS, the Login accounts, database links, Agent Jobs, and key configurations of the source database also need to be synchronized to the destination database.

Login Account

Login account is an instance-level account of Microsoft SQL Server and is used to manage user server and database permissions. Generally, a user has multiple such accounts. After the user is migrated to the RDS SQL Server DB instance, you need to manually create corresponding Login accounts on the DB instance. The following describes how to create a Login account with the same name and

 Copyright © Huawei Technologies Co., Ltd.

password as those of your local Login account on the RDS SQL Server DB instance and grant permissions to the account.

Step 1 Execute the following script to obtain the script for creating a Local account on your local instance. The obtained script can be directly executed on the destination DB instance to create a Login account with the same name and password.

```
SELECT 'IF (SUSER_ID('+QUOTENAME(SP.name,'')+') IS NULL) BEGIN CREATE LOGIN '
+QUOTENAME(SP.name)+
CASE
WHEN SP.type_desc = 'SQL_LOGIN' THEN ' WITH PASSWORD = '
+CONVERT(NVARCHAR(MAX),SL.password_hash,1)+ ' HASHED,SID='
+CONVERT(NVARCHAR(MAX),SP.SID,1)+',CHECK_EXPIRATION = '
+ CASE WHEN SL.is_expiration_checked = 1 THEN 'ON' ELSE 'OFF' END +', CHECK_POLICY = '
+CASE WHEN SL.is_policy_checked = 1 THEN 'ON,' ELSE 'OFF,' END
ELSE ' FROM WINDOWS WITH'
END
+' DEFAULT_DATABASE=[' +SP.default_database_name+ '], DEFAULT_LANGUAGE=['
+SP.default_language_name+ '] END;' as CreateLogin
FROM sys.server_principals AS SP LEFT JOIN sys.sql_logins AS SL
ON SP.principal_id = SL.principal_id
WHERE SP.type ='S'
AND SP.name NOT LIKE '##%##'
AND SP.name NOT LIKE 'NT AUTHORITY%'
AND SP.name NOT LIKE 'NT SERVICE%'
AND SP.name NOT IN ('rdsadmin','rdsbackup','rdsuser','rdsmirror','public')
```

Step 2 Execute the script in **Step 1**:

Figure 3-25 Obtaining the script

	CreateLogin
1	IF (SUSER_ID('sa') IS NULL) BEGIN CREATE LOGIN [sa] WITH PASSWORD = 0x010039BF2EFAD6A3DB4E2AEB941E8BD32E5189A4EE757...
2	IF (SUSER_ID('rdsuser2') IS NULL) BEGIN CREATE LOGIN [rdsuser2] WITH PASSWORD = 0x0100EE8BCBC25FC67008D4EE75AD660D1...
3	IF (SUSER_ID('csidbo') IS NULL) BEGIN CREATE LOGIN [csidbo] WITH PASSWORD = 0x0100A508789C15CE688664B162A5EDF4F4D2E...
4	IF (SUSER_ID('TestLogin7') IS NULL) BEGIN CREATE LOGIN [TestLogin7] WITH PASSWORD = 0x010073DA9A79E6677E8AF7077EF67...
5	IF (SUSER_ID('rdsuser3') IS NULL) BEGIN CREATE LOGIN [rdsuser3] WITH PASSWORD = 0x01009448FEDBCE8D5B5B2529384028CA0...
6	IF (SUSER_ID('Test2') IS NULL) BEGIN CREATE LOGIN [Test2] WITH PASSWORD = 0x0100130953CEEABC997D08B6BAF65F84EBCAA44...
7	IF (SUSER_ID('Test3') IS NULL) BEGIN CREATE LOGIN [Test3] WITH PASSWORD = 0x0100EE98873948E02595BDCD953B426637281B7...
8	IF (SUSER_ID('Test4') IS NULL) BEGIN CREATE LOGIN [Test4] WITH PASSWORD = 0x01000EE91B9BF087741F16A44E70AA813D0BA8B...
9	IF (SUSER_ID('Test5') IS NULL) BEGIN CREATE LOGIN [Test5] WITH PASSWORD = 0x010056EBF845DF098D2DF9395AF7E7618A20735...

Step 3 Copy and execute the script obtain in **Step 2** on the destination instance. The created Login account is the same as the original one.

Step 4 Map the newly created Login account to the database user permissions that have been migrated to the RDS SQL Server DB instance to ensure permission consistency.

```
declare @DBName nvarchar(200)
declare @Login_name nvarchar(200)
declare @SQL nvarchar(MAX)
set @Login_name = 'TestLogin7' //Enter the login name one by one.
declare DBName_Cursor cursor for
select quotename(name)from sys.databases where  database_id > 4 and state = 0
and name not like '%$%'
and name <> 'rdsadmin'
open DBName_Cursor
fetch next from DBName_Cursor into @DBName
WHILE @@FETCH_STATUS= 0
begin
SET @SQL='    USE '+ (@DBName)+ '
if exists(select top 1 1 from sys.sysusers where name = '''+ @Login_name +''')
begin
```

 Copyright © Huawei Technologies Co., Ltd.

```
ALTER USER '+@Login_name+' with login = '+@Login_name+';
end
'
print @SQL
EXEC (@SQL)
fetch next from DBName_Cursor into @DBName
end
close DBName_Cursor
deallocate DBName_Cursor
```

NOTE

After the preceding script is executed, you can view the Login account with the same name on the new instance, and the password and permission are the same as those on your local host.

----End

Database Link

SQL Server allows you to create database links to interact with databases on external DB instances. Therefore you can query, synchronize, and compare databases of different types or on different DB instances. However, these links cannot be automatically synchronized to the DB instance on cloud so you need to synchronize them manually.

Step 1 Connect the local DB instance and cloud DB instance through Microsoft SQL Server Management Studio. Choose **Server Objects** > **Linked Servers** and locate the DBLink of the current DB instance.

Figure 3-26 Viewing database links

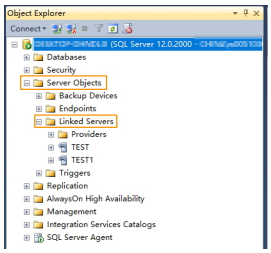

Step 2 Select the linked server and press **F7**. The **Object Explore** page is displayed. On this page, you can quickly create a script.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-27 Creating the script

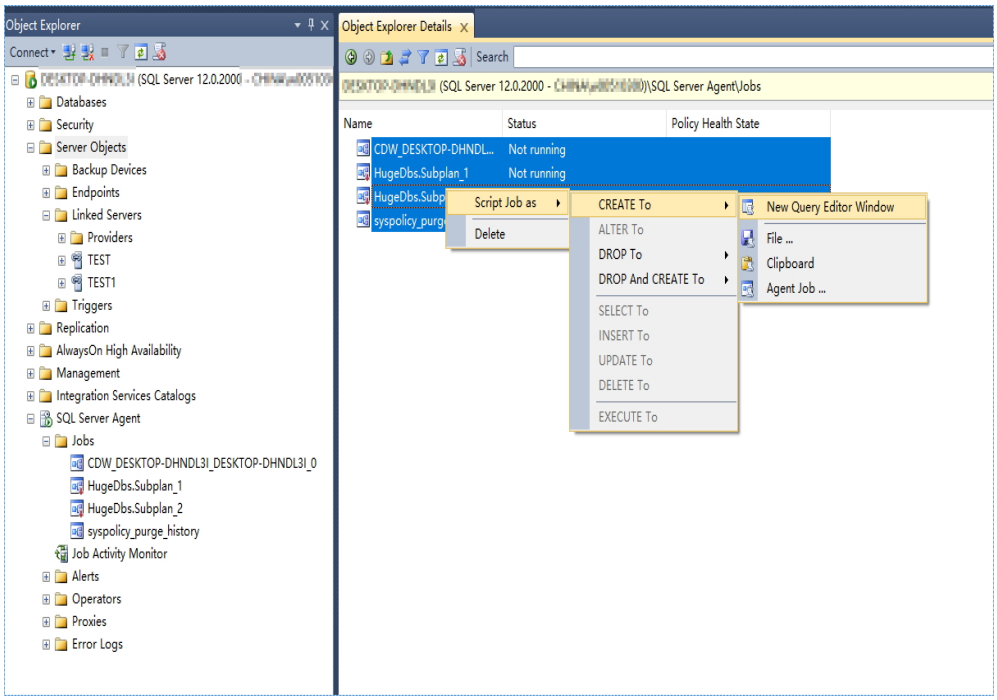

Step 3 In the displayed window, view all the scripts for creating DBLinks of the current DB instance. You only need to copy the scripts to the destination DB instance and change the password on @rmtpassword.

```
USE [master]
GO

/****** Object:  LinkedServer [DRS_TEST_REMOTE]    Script Date: 2019/5/25 17:51:50 ******/
EXEC master.dbo.sp_addlinkedserver @server = N'DRS_TEST_REMOTE', @srvproduct=N'',
@provider=N'SQLNCLI', @datasrc=N'DESKTOP-B18JH5T\SQLSERVER2016EE'
/* For security reasons the linked server remote logins password is changed with ######## */
EXEC master.dbo.sp_addlinkedsrvlogin
@rmtsrvname=N'DRS_TEST_REMOTE',@useself=N'False',@locallogin=NULL,@rmtuser=N'sa',@rmtpassword='########'
GO
```

NOTE

The preceding script is an example. The created script may contain a large number of default system configuration items. You need to retain only the following two key scripts for each DBLink. In addition, you need to enter the account and password again.

----End

Agent JOB

Agent Job is the agent service of Microsoft SQL Server. It helps you quickly create scheduled tasks on DB instances, perform routine O&M, and process data. You need to manually migrate local Job scripts.

Step 1 Connect the local DB instance and cloud DB instance through Microsoft SQL Server Management Studio. Choose **SQL Server Agent** > **Jobs** and locate all the jobs of the current DB instance.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-28 Viewing Jobs

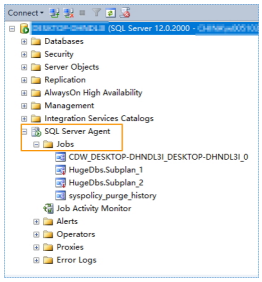

Step 2 Select a job and press **F7**. All jobs are displayed on the **Object Explore** page. Select all jobs and create a script in the new window.

Figure 3-29 Creating a script

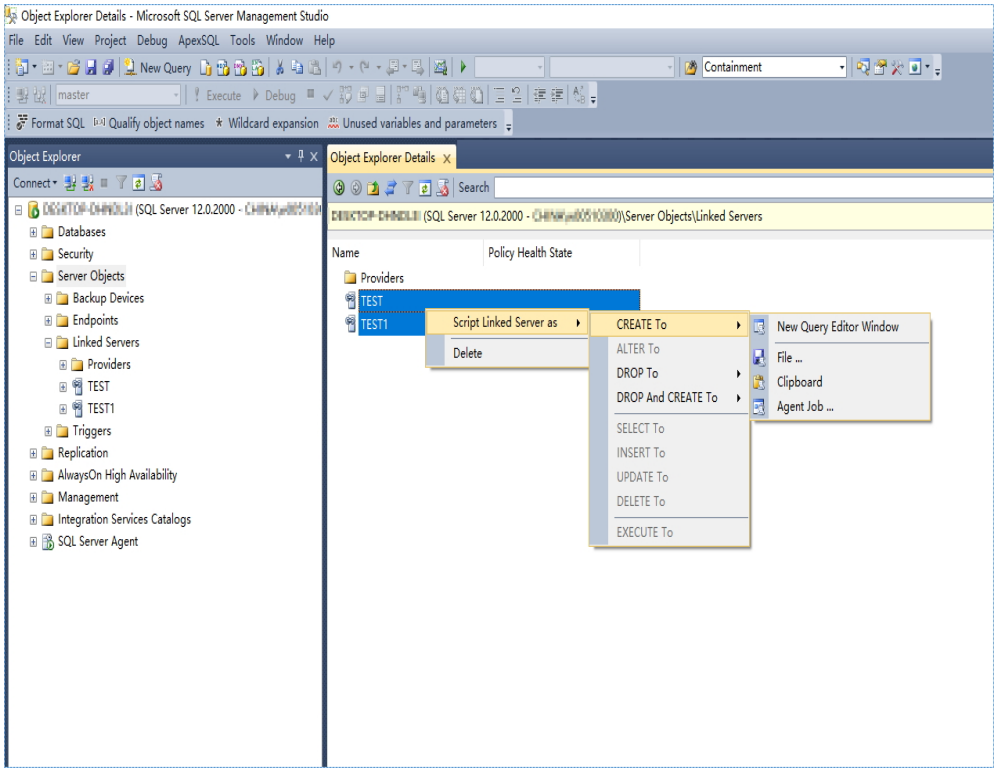

Step 3 Copy the T-SQL script in the new window to the new DB instance, and then modify the following key items to ensure that the creation is successful.

- Modify the owner account of each job.

  Example:

  @owner_login_name=N'rdsuser'

- Modify the DB instance name of each job.

  Example:

  @server=N' DB instance IP address'

  @server_name = N'DB instance IP address'

 Copyright © Huawei Technologies Co., Ltd.

NOTE

The owner account of the new job is very important. On the RDS SQL Server DB instance, only the owner of the job can view the job of the DB instance. Therefore, it is recommended that all job owners use the same account to facilitate job management.

----End

Key Configuration

After the database is restored to the RDS SQL Server DB instance, some local important configuration items need to be synchronized to keep service running properly.

1. tempdb: The file configuration of the temporary database needs to be synchronized.

   It is recommended that you set 8 temporary files and ensure that the files are stored in **D:\RDSDBDATA\Temp**.

   Run the following script on the destination database to add the temporary database file configuration:

```
USE [master]
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb1', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb1.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb2', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb2.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb3', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb3.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb4', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb4.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb5', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb5.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb6', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb6.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb7', FILENAME = N'D:
\RDSDBDATA\Temp\tempdb7.ndf' , SIZE = 65536KB , FILEGROWTH = 65536KB )
GO
```

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-30 Checking temporary files

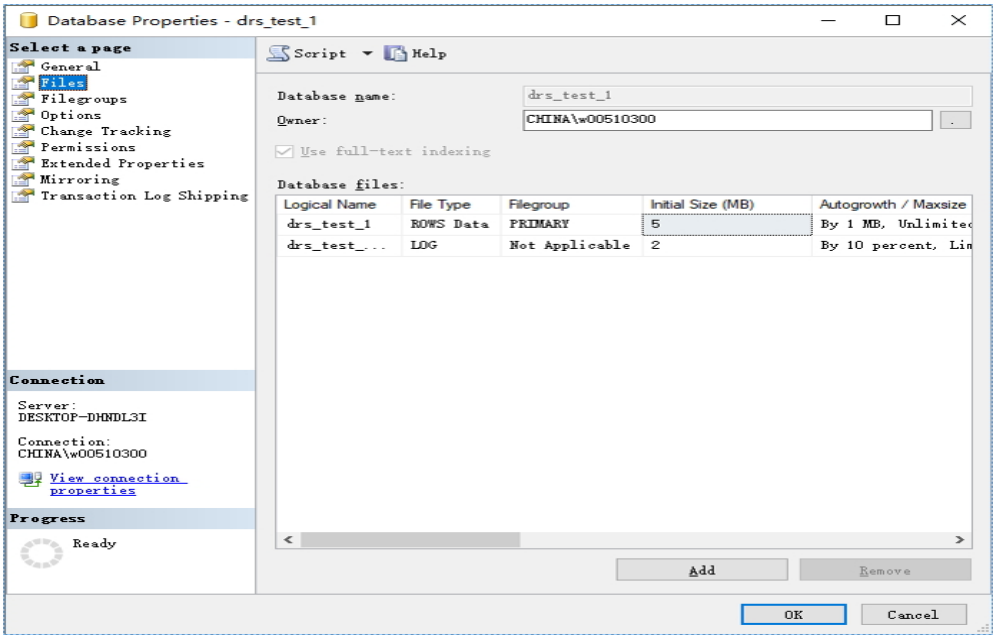

2. Database isolation level: Check whether the database isolation level is enabled on the source DB instance and synchronize the isolation level to the RDS SQL Server DB instance. There are two snapshot isolation parameters:
   - Is Read Committed Snapshot On
   - Allow Snapshot Isolation

   If the database isolation level of the source DB instance is enabled, you can run the following script on the destination database to enable the database isolation level:

   ```
   USE [DBName]
   GO
   ALTER DATABASE [DBName] SET READ_COMMITTED_SNAPSHOT ON WITH NO_WAIT
   GO
   ALTER DATABASE [DBName] SET ALLOW_SNAPSHOT_ISOLATION ON
   GO
   ```

3. Max Degree of Parallelism: The maximum degree of parallelism is set to **0** by default on the RDS SQL Server instance. You can also set the value based on the local settings to avoid exceptions in different service scenarios.

   In **Object Explorer**, right-click a local server and select **Properties**. Click the **Advanced** node. In the **Max Degree of Parallelism** box, view the value of the local instance and change the **max degree of parallelism** value in the parameter group of the destination RDS SQL Server instance to the same.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-31 Max Degree of Parallelism

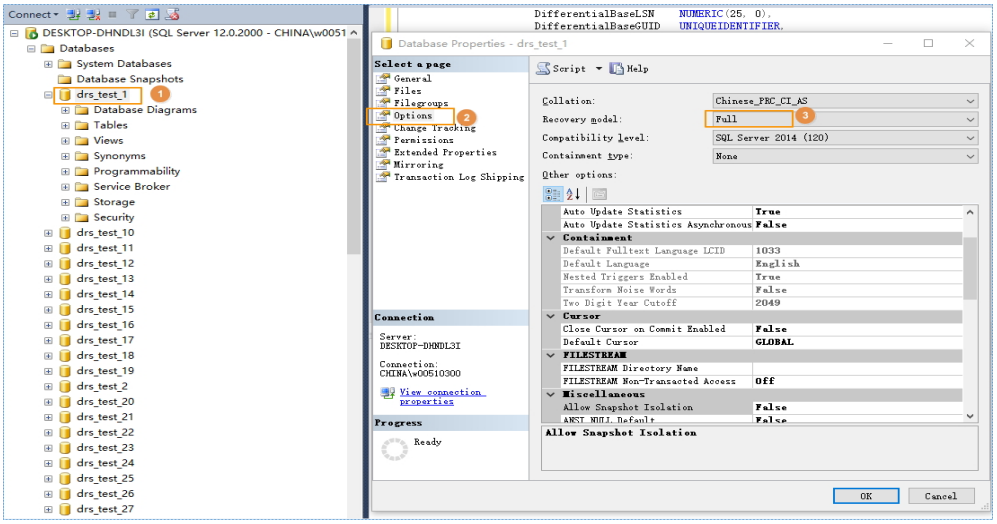

Log in to the RDS console. On the **Instance Management** page, click the target DB instance name. Choose **Parameters**, search for the **max degree of parallelism** parameter, and change its value.

Figure 3-32 max degree of parallelism

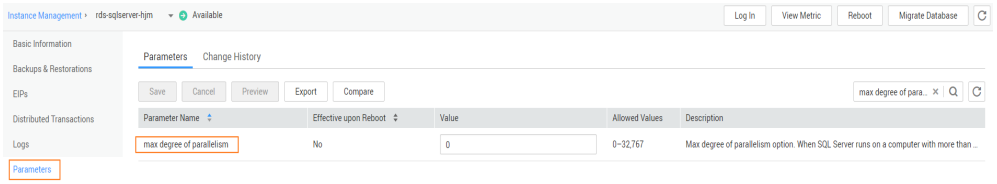

4. Check whether the database recovery model on the cloud is set to **Full**. If not, change the mode.

   Right-click the database and choose **Properties** from the shortcut menu. In the displayed page, select **Options**. Then, verify that **Recovery Model** is set to **Full**. Ensure that the database is highly available and the backup policy is executable.

Figure 3-33 Checking the database recovery model

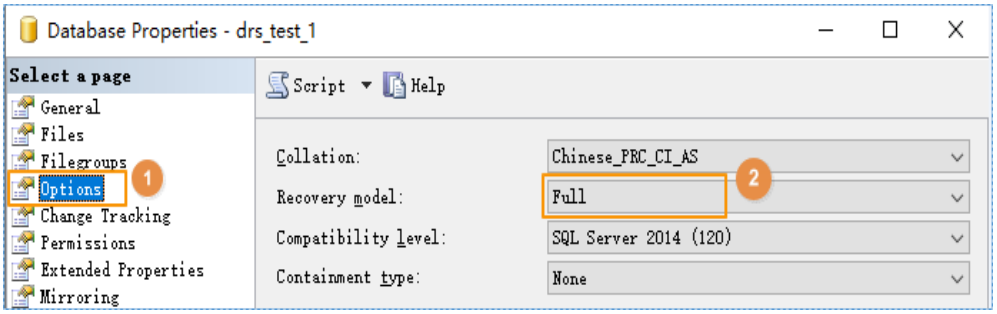

 Copyright © Huawei Technologies Co., Ltd.

4 Real-Time Synchronization

4.1 From Other Cloud PostgreSQL to RDS PostgreSQL

DRS helps you synchronize PostgreSQL instances from other cloud platforms to the current cloud. DRS supports real-time synchronization to ensure real-time flow of key service data.

This section describes how to use DRS to synchronize data from a PostgreSQL instance on another cloud to the current cloud. Synchronization scenarios include:

- Synchronizing PostgreSQL databases from another cloud to the current cloud.
- Synchronizing PostgreSQL databases from other cloud servers to the current cloud in real time.

Diagram

Figure 4-1 Real-time synchronization of other cloud RDS PostgreSQL databases

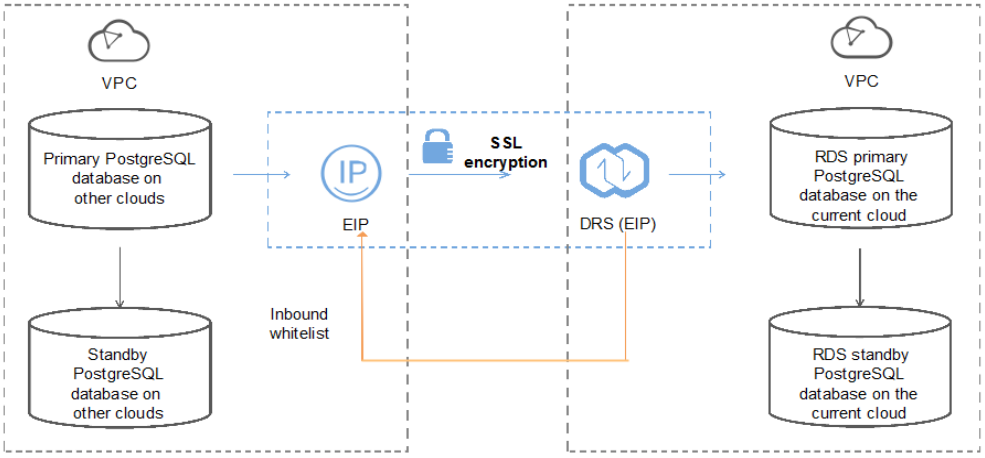

 Copyright © Huawei Technologies Co., Ltd.

Figure 4-2 PostgreSQL databases on other cloud servers

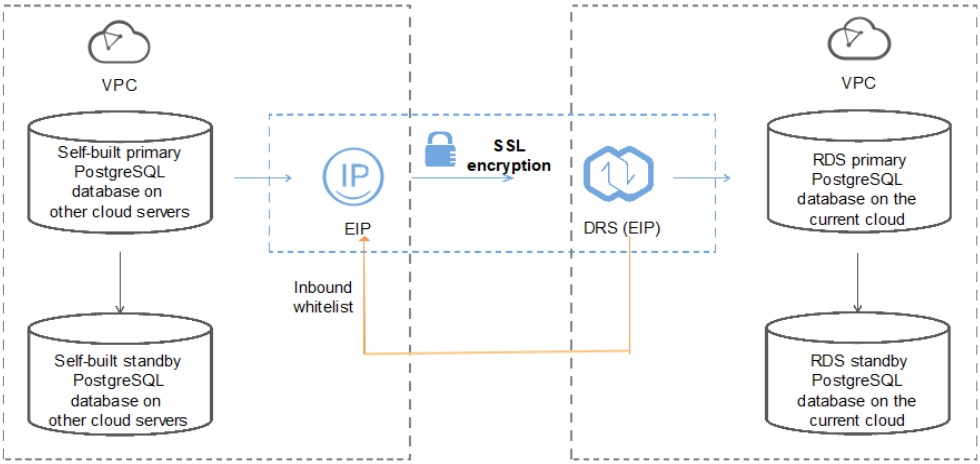

Synchronization Process

Figure 4-3 Flowchart

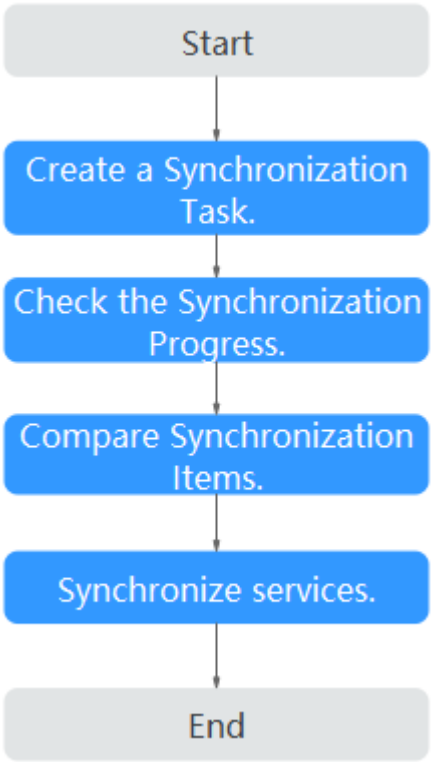

Synchronization Suggestions (Important)

- Database synchronization is closely impacted by a wide range of environmental and operational factors. To ensure the synchronization goes smoothly, you are advised to perform a test run before the actual synchronization to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.

 Copyright © Huawei Technologies Co., Ltd.

- It is strongly recommended that you start a task during off-peak hours. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

Notes on Synchronization (Important)

NOTICE

Before creating a real-time synchronization task, read this section carefully.

For details, see **Precautions** in *Real-Time Synchronization*.

Synchronization Preparations

1. Permissions

   Table 4-1 lists the permissions required for the source and destination databases when you synchronize a PostgreSQL database from another cloud to the current cloud.

   Table 4-1 Account permissions

Synchronization Type	Full	Full+Incremental
Source	• Database: CONNENT • Schema: USAGE • Table with a primary key: SELECT • Tables without primary keys: SELECT, UPDATE, DELETE, and TRUNCATE • Sequence: SELECT	• Account: REPLICATION • Database: CONNENT • Schema: USAGE • Table with a primary key: SELECT • Tables without primary keys: SELECT, UPDATE, DELETE, and TRUNCATE • Sequence: SELECT
Destination	CREATEDB	CREATEDB, CREATEROLE

   - Source database permissions:

     The source PostgreSQL database account must have all the required permissions listed in **Table 4-1**. If the permissions are insufficient, create a user that has all of the permissions on the source database.

     Currently, you can use DRS to migrate databases on Alibaba Cloud and Tencent Cloud.

 Copyright © Huawei Technologies Co., Ltd.

- Destination database permissions:

  If the destination database is a PostgreSQL database on the current cloud, the initial account can be used.

2. Network settings

   Enable public accessibility for the source database.

   - Source database network settings:

     Enable public accessibility for the source PostgreSQL database.

   - Destination database network settings:

     By default, the destination database and the DRS replication instance are in the same VPC and can communicate with each other. No further configuration is required.

3. Security rules

   - Source database security group settings:

     The synchronization instance needs to be able to access the source database, which means that the EIP of the replication instance must be on the whitelist of the source PostgreSQL instance.

     For example, if you want to synchronize a PostgreSQL database from Alibaba Cloud to the current cloud, you need to set the source database whitelist based on the instructions provided by **Alibaba Cloud documentation**. Before configuring the network whitelist, you need to obtain the EIP of the synchronization instance.

     After creating a synchronization instance on the DRS console, you can find the EIP on the **Configure Source and Destination Databases** page.

     Figure 4-4 Synchronization Instance EIP

     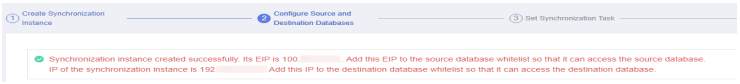

     You can also add 0.0.0.0/0 to the source database whitelist to allow any IP address to access the source database but this action may result in security risks.

     After the synchronization is complete, you can delete the rules.

   - Destination database security group settings:

     By default, the destination database and the DRS synchronization instance are in the same VPC and can communicate with each other. No further configuration is required.

4. Other

   DRS can synchronize only some types of DDL statements of PostgreSQL. For details, see **related guides**. Properly plan the source database services. Do not perform DDL operations that are not supported during synchronization.

Procedure

Step 1 Create a synchronization task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.

 Copyright © Huawei Technologies Co., Ltd.

2. On the **Data Synchronization Management** page, click **Create Synchronization Task**.
3. On the displayed page, specify the task name, description, and synchronization instance details.

Figure 4-5 Synchronization instance information

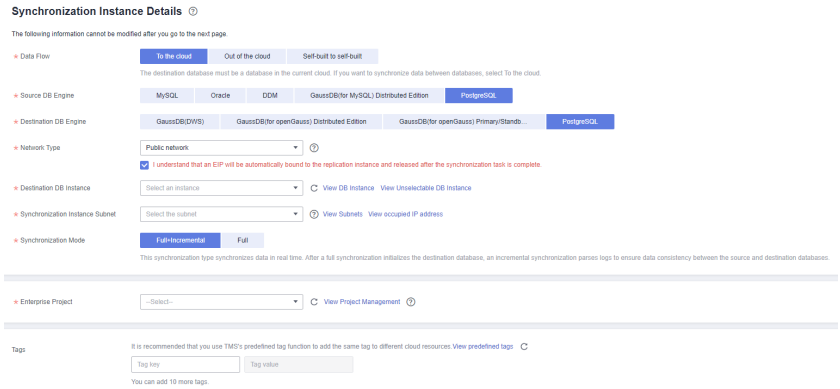

Table 4-2 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.
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: =<>&'\"

Table 4-3 Synchronization instance settings

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	Select **PostgreSQL**.
Destination DB Engine	Select **PostgreSQL**.
Network Type	Select **Public network**. Enabling SSL is recommended. It may slow down the synchronization by 20% to 30% but it ensures data security.
Destination DB Instance	The PostgreSQL instance you purchased.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Synchronization Type	**Full+Incremental** This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. **NOTE** If you select the **Full+Incremental** synchronization mode, ongoing changes made to the data will be synchronized to the destination database in real time, ensuring that the source database remain accessible.

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

Figure 4-6 Source and destination database details

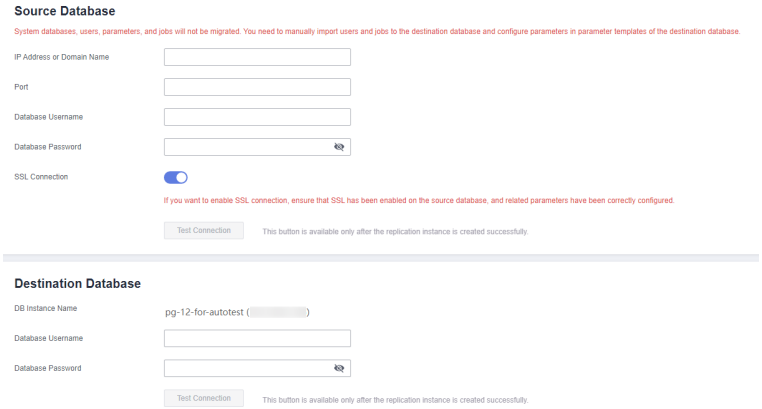

Table 4-4 Source database settings

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	A username for the source database.
Database Password	The password for the database username.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
SSL Connection	To improve data security during synchronization over a public network, you are advised to enable SSL to encrypt synchronization links and upload a CA certificate.

Table 4-5 Destination database settings

Parameter	Description
DB Instance Name	The RDS PostgreSQL instance you have selected during the synchronization instance creation is displayed by default and cannot be changed.
Database Username	The username for accessing the destination PostgreSQL DB instance.
Database Password	The password for the database username.

5. On the **Set Synchronization Task** page, select the synchronization object and user.

Figure 4-7 Synchronization mode

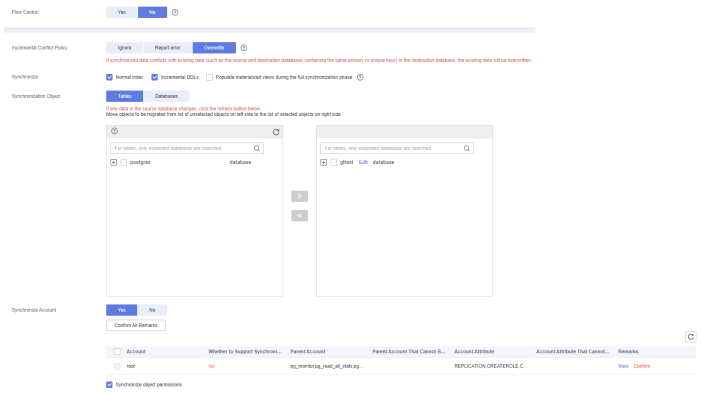

Table 4-6 Synchronization object

Parameter	Description
Flow Control	You can choose whether to control the flow.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Incremental Conflict Policy	The real-time synchronization function provides conflict policies for you to choose from if the synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database. Select any of the following conflict policies: – **Ignore** The system will skip the conflicting data and continue the subsequent synchronization process. – **Report error** The synchronization task will be stopped and fail. – **Overwrite** Conflicting data will be overwritten. In the following scenarios, you can select **Ignore** or **Overwrite**. In other scenarios, you are advised to select **Report error**. – Data already exists in the destination database. – Multiple source databases are synchronized to one destination database. – Data in the destination database is updated manually.
Synchronize	Available options: **Index**, **Incremental DDLs**, and **Populate materialized views during the full synchronization phase** **Populate materialized views during the full synchronization phase**: This option takes effect only for materialized views that was populated in the source database. This operation affects the full synchronization performance. You perform this operation after the full synchronization is complete.
Synchronization Object	DRS supports database- and table-level synchronization. You can select databases or tables for synchronization based on your service requirements. A single task can synchronize objects from only one database. Database name mapping is supported. – Database-level synchronization synchronizes all objects in the source database to the destination database. – Table-level synchronization synchronizes the selected table objects to the destination database. **NOTE** If you select table-level synchronization, the selected table may have dependencies on other objects in the database. If the referenced objects are not selected and do not exist in the destination database, the task will fail. Ensure that all referenced objected are selected before synchronization or select database-level synchronization.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Synchronize Account	During the synchronization, you can synchronize accounts based on your service requirements. There are two types of accounts: accounts that can be synchronized and accounts that cannot be synchronized. For accounts that cannot be synchronized, you can view details in the **Remarks** column and determine whether to synchronize accounts and permissions based on your service requirements.

6. On the **Check Task** page, check the synchronization task.
   - If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

     For details about how to handle check failures, see **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.
   - If all check items are successful, click **Next**.

     NOTE

     You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.
7. On the **Confirm Task** page, specify **Start Time** and click **Next**.

   NOTE

   - Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements.
   - After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
8. After the task is submitted, go back to the **Data Synchronization Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full synchronization

  Viewing the synchronization progress: Click the target full synchronization task, and on the **Migration Progress** tab, you can see the synchronization progress of the structure, data, indexes, and synchronization objects. When the progress reaches 100%, the synchronization is complete.
- Incremental synchronization
  - Viewing the synchronization delay: After the full synchronization is complete, an incremental synchronization starts. On the **Data Synchronization Management** page, click the target synchronization task. On the displayed page, click **Synchronization Progress** to view the synchronization delay of the incremental synchronization. If the synchronization delay is 0s, the destination database is being

 Copyright © Huawei Technologies Co., Ltd.

synchronized with the source database in real time. You can also view the data consistency on the **Synchronization Comparison** tab.

Figure 4-8 Viewing the synchronization delay

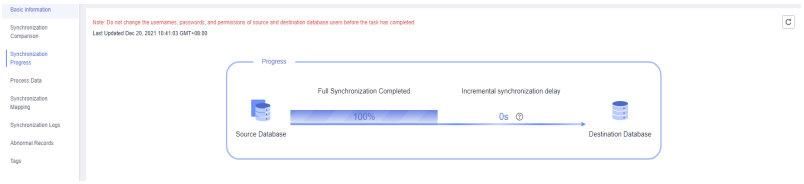

- Viewing the synchronization comparison results: To minimize service downtime, click the name of an incremental synchronization task. On the **Synchronization Comparison** page, create a comparison task.

  For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. 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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.
2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.

   ```
   select * from pg_stat_activity;
   ```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS synchronization instance. If no additional session executes SQL statements, the service has been stopped.
3. On the **Synchronization Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the synchronization is complete.

Step 4 Complete the synchronization.

1. Stop the synchronization task. After databases and services are synchronized to the destination database, to prevent operations on the source database

 Copyright © Huawei Technologies Co., Ltd.

from being synchronized to the destination database to overwrite data, you can stop the synchronization task. This operation only deletes the synchronization instance, and the synchronization task is still displayed in the task list. You can view or delete the task. You will not be charged for the synchronization task after the task is stopped.

2. Delete the synchronization task. After the synchronization task is complete, you can delete it. After the synchronization task is deleted, it will not be displayed in the task list.

----End

4.2 From PostgreSQL on ECS to RDS PostgreSQL

DRS helps you synchronize data from PostgreSQL databases on ECSs to PostgreSQL instances on the current cloud. With DRS, you can synchronize databases online with zero downtime and your services and databases can remain operational during migration.

This section describes how to use DRS to synchronize data from a PostgreSQL database on an ECS to a PostgreSQL instance on the current cloud. The following network scenarios are supported:

- Source and destination databases are in the same VPC.
- Source and destination databases are in different VPCs.

Diagram

Figure 4-9 Source and destination databases in the same VPC

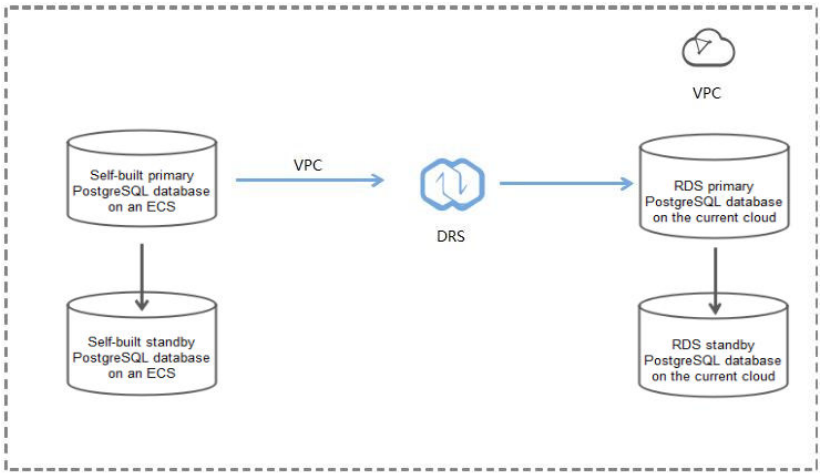

 Copyright © Huawei Technologies Co., Ltd.

Figure 4-10 Source and destination databases in the same region and different VPCs

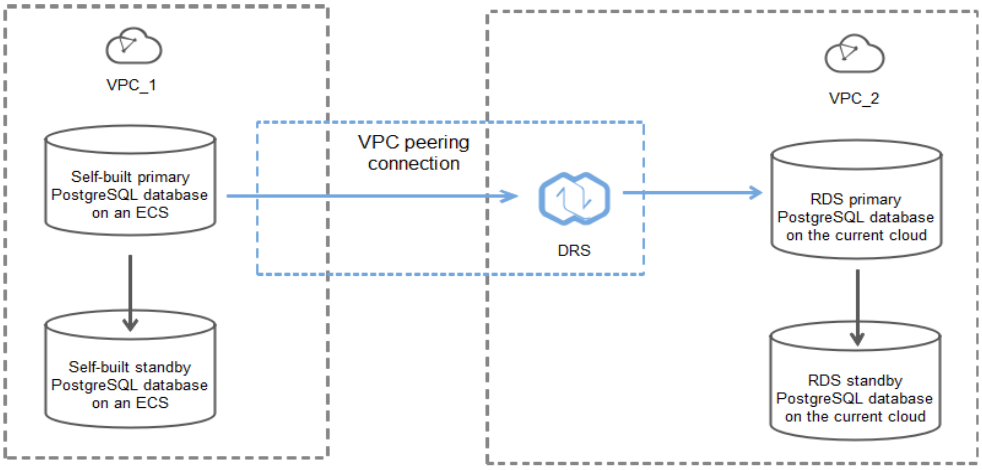

Synchronization Process

Figure 4-11 Flowchart

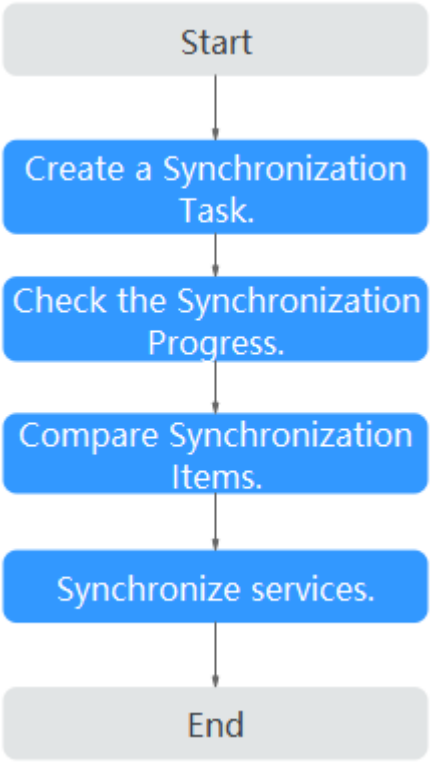

Synchronization Suggestions (Important)

- Database synchronization is closely impacted by a wide range of environmental and operational factors. To ensure the synchronization goes smoothly, you are advised to perform a test run before the actual synchronization to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.

 Copyright © Huawei Technologies Co., Ltd.

- It is strongly recommended that you start a task during off-peak hours. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

Notes on Synchronization (Important)

NOTICE

Before creating a real-time synchronization task, read this section carefully.

For details, see **Precautions** in *Real-Time Synchronization*.

Synchronization Preparations

1. Permissions

   Table 4-7 lists the source and destination database user permissions required in full and incremental synchronization from PostgreSQL databases on ECS to the PostgreSQL DB instances on the current cloud.

   Table 4-7 Account permissions

Synchronization Type	Full	Full+Incremental
Source	• Database: CONNENT • Schema: USAGE • Table with a primary key: SELECT • Tables without primary keys: SELECT, UPDATE, DELETE, and TRUNCATE • Sequence: SELECT	• Account: REPLICATION • Database: CONNENT • Schema: USAGE • Table with a primary key: SELECT • Tables without primary keys: SELECT, UPDATE, DELETE, and TRUNCATE • Sequence: SELECT
Destination	CREATEDB	CREATEDB, CREATEROLE

   - Source database permissions:

     The source database user must have all the required permissions listed in **Table 4-7**. If the permissions are insufficient, create a user that has all of the permissions on the source database.
   - Destination database permissions:

     If the destination database is a PostgreSQL database on the current cloud, the initial account can be used.

 Copyright © Huawei Technologies Co., Ltd.

2. Network settings
   - The source and destination databases must be in the same region.
   - The source and destination databases can be either in the same VPC or different VPCs.
     - If the source and destination databases are in different VPCs, the subnets of the source and destination databases are required to be in different CIDR blocks. You need to create a VPC peering connection between the two VPCs. For details, see .
     - If the source and destination databases are in the same VPC, the networks are interconnected by default.
3. Security rules
   - In the same VPC, the network is connected by default. You do not need to set a security group.
   - In different VPCs, establish a VPC peering connection between the two VPCs. You do not need to set a security group.
4. Other

   DRS can synchronize only some types of DDL statements of PostgreSQL. For details, see **related guides**. Properly plan the source database services. Do not perform DDL operations that are not supported during synchronization.

Procedure

Step 1 Create a synchronization task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.
2. On the **Data Synchronization Management** page, click **Create Synchronization Task**.
3. On the displayed page, specify the task name, exception notification, task description, and synchronization instance details.

Figure 4-12 Synchronization instance settings

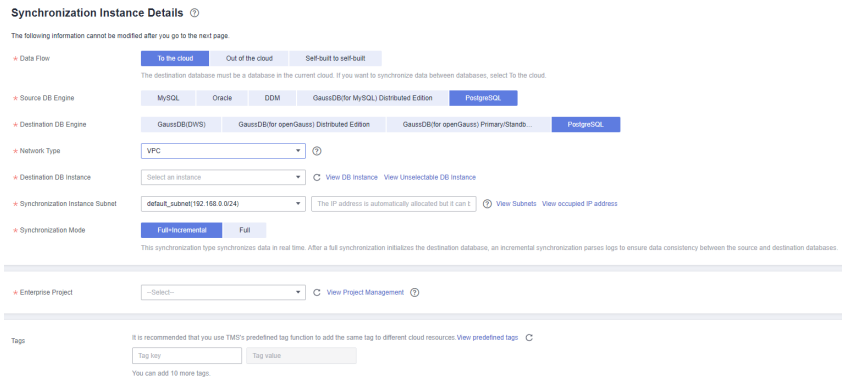

 Copyright © Huawei Technologies Co., Ltd.

Table 4-8 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.
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: =<>&'\"

Table 4-9 Synchronization instance settings

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	Select **PostgreSQL**.
Destination DB Engine	Select **PostgreSQL**.
Network Type	Select a VPC network. Enabling SSL is recommended. It may slow down the synchronization by 20% to 30% but it ensures data security.
Destination DB Instance	The PostgreSQL instance you purchased.
Synchronization Type	**Full+Incremental** This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. **NOTE** If you select the **Full+Incremental** synchronization type, data generated during the full synchronization will be synchronized to the destination database with zero downtime, ensuring that both the source and destination databases remain accessible.

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

 Copyright © Huawei Technologies Co., Ltd.

NOTE

Source databases are classified into two types: self-built databases on ECSs and RDS DB instances. Configure parameters based on different scenarios.

- Scenario 1: Self-built databases on ECS - source database configuration

Figure 4-13 Self-build on ECS - source database information

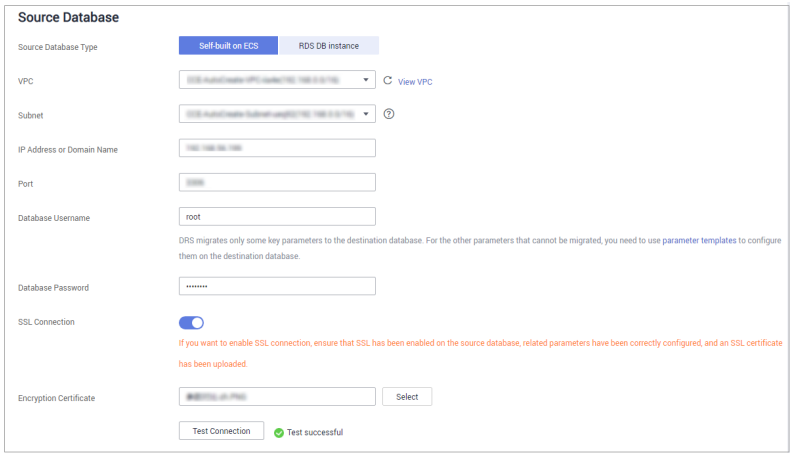

Table 4-10 Self-build on ECS - source database information

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 logically 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	Enter an integer ranging from 1 to 65535, which indicates the port number of the source database.
Database Username	A username for the source database.
Database Password	The password for the database username.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
SSL Connection	To improve data security during synchronization over a public network, you are advised to enable SSL to encrypt synchronization links and upload a CA certificate.

- Scenario 2: RDS DB instance - source database configuration

Figure 4-14 RDS DB instance - source database information

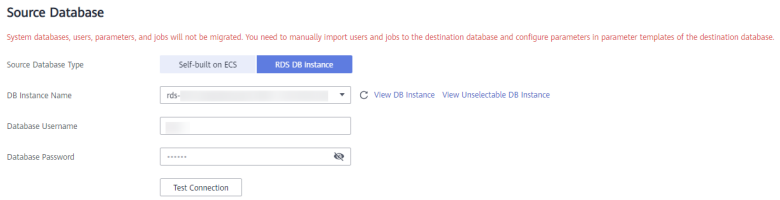

Table 4-11 RDS DB instance - source database information

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

Figure 4-15 Destination database settings

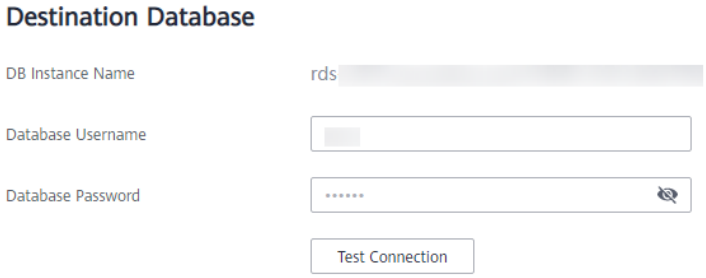

Copyright © Huawei Technologies Co., Ltd.

Table 4-12 Destination database settings

Parameter	Description
DB Instance Name	The RDS PostgreSQL instance you have selected during the synchronization instance creation is displayed by default and cannot be changed.
Database Username	The username for accessing the destination PostgreSQL DB instance.
Database Password	The password for the database username.

5. On the **Set Synchronization Task** page, select the synchronization object and user.

Figure 4-16 Synchronization mode

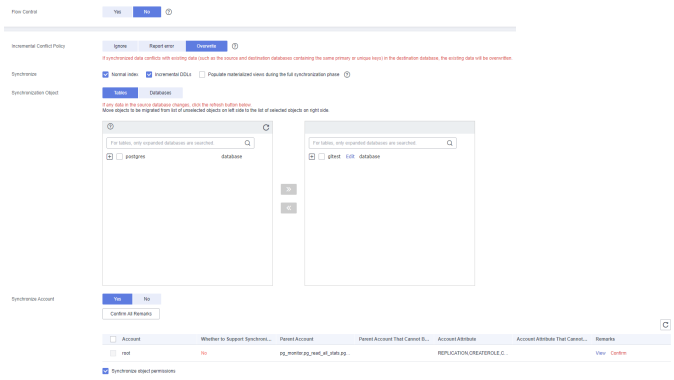

 Copyright © Huawei Technologies Co., Ltd.

Table 4-13 Synchronization object

Parameter	Description
Incremental Conflict Policy	The real-time synchronization function provides conflict policies for you to choose from if the synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database. Select any of the following conflict policies: – **Ignore** The system will skip the conflicting data and continue the subsequent synchronization process. – **Report error** The synchronization task will be stopped and fail. – **Overwrite** Conflicting data will be overwritten. In the following scenarios, you can select **Ignore** or **Overwrite**. In other scenarios, you are advised to select **Report error**. – Data already exists in the destination database. – Multiple source databases are synchronized to one destination database. – Data in the destination database is updated manually.
Synchronize	Available options: **Index**, **Incremental DDLs**, and **Populate materialized views during the full synchronization phase** **Populate materialized views during the full synchronization phase**: This option takes effect only for materialized views that was populated in the source database. This operation affects the full synchronization performance. You perform this operation after the full synchronization is complete.
Synchronization Object	Objects can be synchronized at the database or table level. You can synchronize the entire database or some tables based on service requirements. For a single task, you can synchronize objects in only one database of an instance. Database name mapping is supported. – Database-level synchronization synchronizes all objects in the source database to the destination database. – Table-level synchronization synchronizes the selected table objects to the destination database. **NOTE** If you select table-level synchronization, the selected table may have dependencies on other objects in the database. If the referenced objects are not selected and do not exist in the destination database, the task will fail. Ensure that all referenced objected are selected before synchronization or select database-level synchronization.

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Synchronize Account	During the synchronization, you can synchronize accounts based on your service requirements. There are two types of accounts: accounts that can be synchronized and accounts that cannot be synchronized. For accounts that cannot be synchronized, you can view details in the **Remarks** column and determine whether to synchronize accounts and permissions based on your service requirements.

6. On the **Check Task** page, check the synchronization task.
   - If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

     For details about how to handle check failures, see **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.
   - If all check items are successful, click **Next**.

     NOTE

     You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.
7. On the **Confirm Task** page, specify **Start Time** and click **Next**.

   NOTE

   - Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements.
   - After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
8. After the task is submitted, go back to the **Data Synchronization Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full synchronization

  Viewing the synchronization progress: Click the target full synchronization task, and on the **Migration Progress** tab, you can see the synchronization progress of the structure, data, indexes, and synchronization objects. When the progress reaches 100%, the synchronization is complete.
- Incremental synchronization
  - Viewing the synchronization delay: After the full synchronization is complete, an incremental synchronization starts. On the **Data Synchronization Management** page, click the target synchronization task. On the displayed page, click **Synchronization Progress** to view the synchronization delay of the incremental synchronization. If the synchronization delay is 0s, the destination database is being

 Copyright © Huawei Technologies Co., Ltd.

synchronized with the source database in real time. You can also view the data consistency on the **Synchronization Comparison** tab.

Figure 4-17 Viewing the synchronization delay

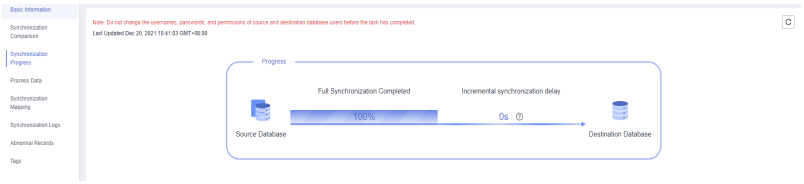

- Viewing the synchronization comparison results: To minimize service downtime, click the name of an incremental synchronization task. On the **Synchronization Comparison** page, create a comparison task.

  For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. 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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.
2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.

   ```
   select * from pg_stat_activity;
   ```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS synchronization instance. If no additional session executes SQL statements, the service has been stopped.
3. On the **Synchronization Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the synchronization is complete.

Step 4 Complete the synchronization.

1. Stop the synchronization task. After databases and services are synchronized to the destination database, to prevent operations on the source database

 Copyright © Huawei Technologies Co., Ltd.

from being synchronized to the destination database to overwrite data, you can stop the synchronization task. This operation only deletes the synchronization instance, and the synchronization task is still displayed in the task list. You can view or delete the task. You will not be charged for the synchronization task after the task is stopped.

2. Delete the synchronization task. After the synchronization task is complete, you can delete it. After the synchronization task is deleted, it will not be displayed in the task list.

----End

4.3 From On-Premises PostgreSQL to RDS PostgreSQL

DRS supports data synchronization from on-premises MySQL databases to PostgreSQL instances on the current cloud. With DRS, you can synchronize databases online with zero downtime and your services and databases can remain operational during migration.

This section describes how to configure DRS to migrate data from an on-premises PostgreSQL database to a PostgreSQL instance on the current cloud. The following network types are supported:

- VPN
- Public network

Diagram

Figure 4-18 VPN network

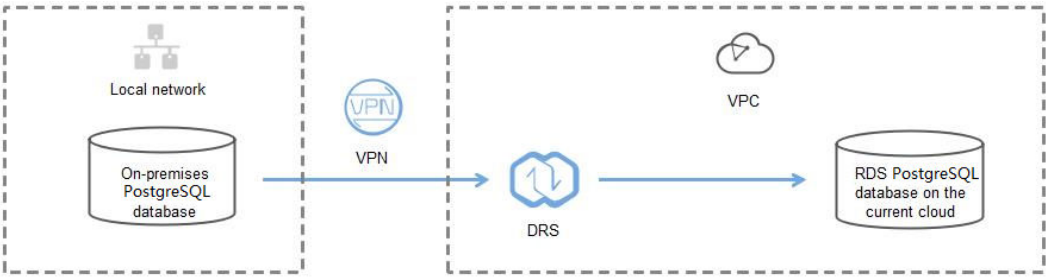

Figure 4-19 Public network+SSL connection

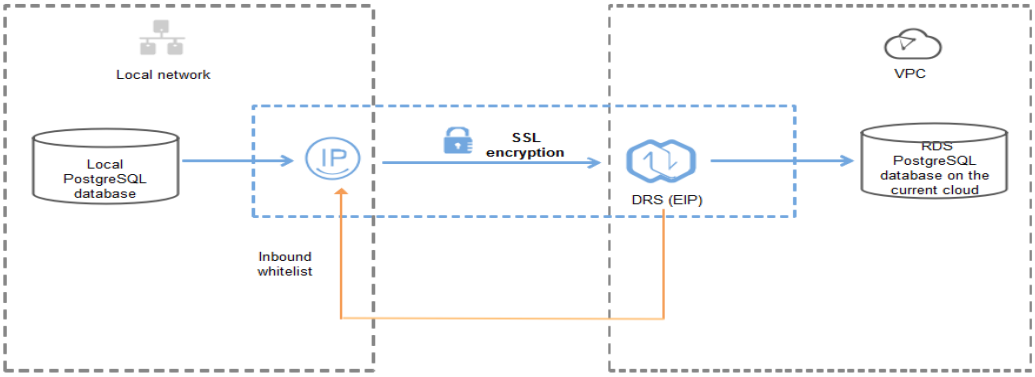

 Copyright © Huawei Technologies Co., Ltd.

Synchronization Process

Figure 4-20 Flowchart

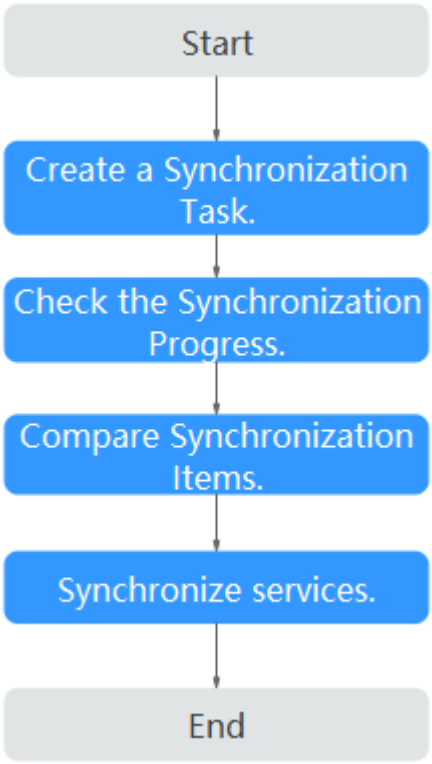

Synchronization Suggestions (Important)

- Database synchronization is closely impacted by a wide range of environmental and operational factors. To ensure the synchronization goes smoothly, you are advised to perform a test run before the actual synchronization to help you detect and resolve any potential issues in advance. Recommendations on how to minimize any potential impacts on your data base are provided in this section.
- It is strongly recommended that you start a task during off-peak hours. A less active database is easier to synchronize successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the synchronization.

Notes on Synchronization (Important)

NOTICE

Before creating a real-time synchronization task, read this section carefully.

For details, see **Precautions** in *Real-Time Synchronization*.

Copyright © Huawei Technologies Co., Ltd.

Synchronization Preparations

1. Permissions

   Table 4-14 lists the source and destination database user permissions required in full and incremental synchronizations from on-premises PostgreSQL databases to PostgreSQL DB instances on the current cloud.

   Table 4-14 Account permissions

Synchronization Type	Full	Full+Incremental
Source	• Database: CONNENT • Schema: USAGE • Table with a primary key: SELECT • Tables without primary keys: SELECT, UPDATE, DELETE, and TRUNCATE • Sequence: SELECT	• Account: REPLICATION • Database: CONNENT • Schema: USAGE • Table with a primary key: SELECT • Tables without primary keys: SELECT, UPDATE, DELETE, and TRUNCATE • Sequence: SELECT
Destination	CREATEDB	CREATEDB, CREATEROLE

   - Source database permissions:

     The source database user must have all the required permissions listed in Table 4-14. If the permissions are insufficient, create a user that has all of the permissions on the source database.
   - Destination database permissions:

     If the destination database is a PostgreSQL database on the current cloud, the initial account can be used.

2. Network settings
   - Source database network settings:

     You can synchronize data from on-premises PostgreSQL databases to RDS PostgreSQL databases on the current cloud through a VPN or public network. Enable public accessibility or establish a VPN for the on-premises PostgreSQL databases based on your service requirements. You are advised to synchronize data through a public network, which is more convenient and cost-effective.
   - Destination database network settings:
     - If you want to access the destination databases through a VPN, enable the VPN service first so that the source database can communicate with the destination database.

 Copyright © Huawei Technologies Co., Ltd.

- If the source database attempts to access the destination database through a public network, no further configuration is required.

3. Security rules
   a. Source database security group settings:
      - If the synchronization is performed over a public network, add the EIP of the DRS synchronization instance to the network whitelist of the source PostgreSQL database to enable the source database to communicate with the current cloud. Before configuring the network whitelist, you need to obtain the EIP of the synchronization instance.

        The IP address on the **Configure Source and Destination Databases** page is the EIP of the synchronization instance.

        Figure 4-21 Synchronization instance EIP

        

      - If the synchronization is performed over a VPN network, add the private IP address of the DRS synchronization instance to the network whitelist of the source database to enable the source database to communicate with the destination database.

        The IP address on the **Configure Source and Destination Databases** page is the private IP address of the synchronization instance.

      If you do take this step, then once the synchronization is complete, you should delete this item from the whitelist or your system will insecure.

   b. Destination database security group settings:

      By default, the destination database and the DRS synchronization instance are in the same VPC and can communicate with each other. No further configuration is required.

4. Other

   DRS can synchronize only some types of DDL statements of PostgreSQL. For details, see **related guides**. Properly plan the source database services. Do not perform DDL operations that are not supported during synchronization.

Procedure

Step 1 Create a synchronization task.

1. Log in to the management console and choose **Databases** > **Data Replication Service** to go to the DRS console.
2. On the **Data Synchronization Management** page, click **Create Synchronization Task**.
3. On the displayed page, specify the task name, description, and synchronization instance details.

 Copyright © Huawei Technologies Co., Ltd.

Figure 4-22 Synchronization instance information

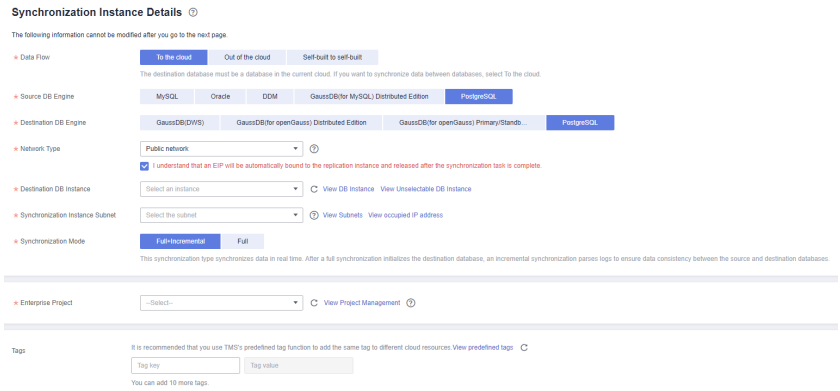

Table 4-15 Task settings

Parameter	Description
Region	The region where your service is running. You can change the region.
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: =<>&'\"

Table 4-16 Synchronization instance settings

Parameter	Description
Data Flow	**To the cloud**
Source DB Engine	Select **PostgreSQL**.
Destination DB Engine	Select **PostgreSQL**.
Network Type	Select **Public network**. Enabling SSL is recommended. It may slow down the synchronization by 20% to 30% but it ensures data security.
Destination DB Instance	The PostgreSQL instance you purchased.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Synchronization Type	**Full+Incremental** This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. **NOTE** If you select the **Full+Incremental** synchronization mode, ongoing changes made to the data will be synchronized to the destination database in real time, ensuring that the source database remain accessible.

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

Figure 4-23 Source and destination database details

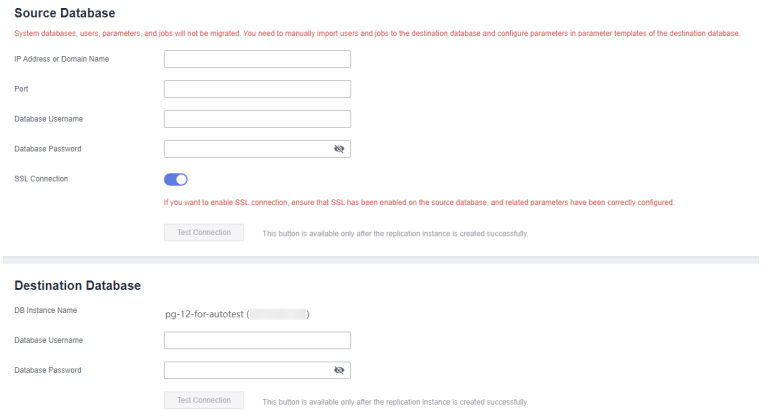

Table 4-17 Source database settings

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

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
SSL Connection	To improve data security during synchronization over a public network, you are advised to enable SSL to encrypt synchronization links and upload a CA certificate.

Table 4-18 Destination database settings

Parameter	Description
DB Instance Name	The RDS PostgreSQL instance you have selected during the synchronization instance creation is displayed by default and cannot be changed.
Database Username	The username for accessing the destination PostgreSQL DB instance.
Database Password	The password for the database username.

5. On the **Set Synchronization Task** page, select the synchronization object and user.

Figure 4-24 Synchronization mode

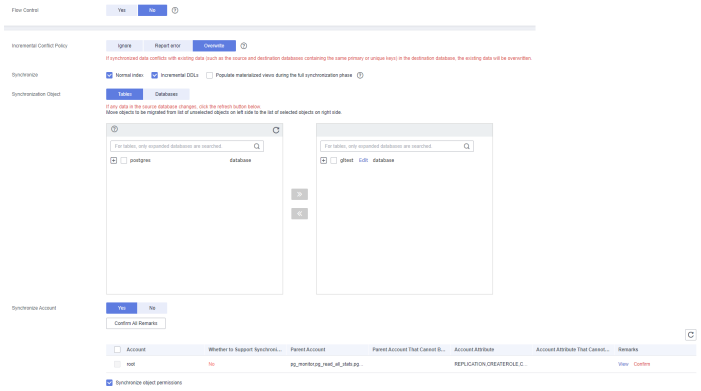

Table 4-19 Synchronization object

Parameter	Description
Flow Control	You can choose whether to control the flow.

 Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Incremental Conflict Policy	The real-time synchronization function provides conflict policies for you to choose from if the synchronized data conflicts with existing data (such as the source and destination databases containing the same primary or unique keys) in the destination database. Select any of the following conflict policies: – **Ignore** The system will skip the conflicting data and continue the subsequent synchronization process. – **Report error** The synchronization task will be stopped and fail. – **Overwrite** Conflicting data will be overwritten. In the following scenarios, you can select **Ignore** or **Overwrite**. In other scenarios, you are advised to select **Report error**. – Data already exists in the destination database. – Multiple source databases are synchronized to one destination database. – Data in the destination database is updated manually.
Synchronize	Available options: **Index**, **Incremental DDLs**, and **Populate materialized views during the full synchronization phase** **Populate materialized views during the full synchronization phase**: This option takes effect only for materialized views that was populated in the source database. This operation affects the full synchronization performance. You perform this operation after the full synchronization is complete.
Synchronization Object	DRS supports database- and table-level synchronization. You can select databases or tables for synchronization based on your service requirements. A single task can synchronize objects from only one database. Database name mapping is supported. – Database-level synchronization synchronizes all objects in the source database to the destination database. – Table-level synchronization synchronizes the selected table objects to the destination database. **NOTE** If you select table-level synchronization, the selected table may have dependencies on other objects in the database. If the referenced objects are not selected and do not exist in the destination database, the task will fail. Ensure that all referenced objected are selected before synchronization or select database-level synchronization.

Copyright © Huawei Technologies Co., Ltd.

Parameter	Description
Synchronize Account	During the synchronization, you can synchronize accounts based on your service requirements. There are two types of accounts: accounts that can be synchronized and accounts that cannot be synchronized. For accounts that cannot be synchronized, you can view details in the **Remarks** column and determine whether to synchronize accounts and permissions based on your service requirements.

6. On the **Check Task** page, check the synchronization task.
   - If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

     For details about how to handle check failures, see **Checking Whether the Source Database Is Connected** in *Data Replication Service User Guide*.
   - If all check items are successful, click **Next**.

     NOTE

     You can proceed to the next step only when all checks are successful. If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.
7. On the **Confirm Task** page, specify **Start Time** and click **Next**.

   NOTE

   - Set **Start Time** to **Start upon task creation** or **Start at a specified time** based on site requirements.
   - After a synchronization task is started, the performance of the source and destination databases may be affected. You are advised to start a synchronization task during off-peak hours.
8. After the task is submitted, go back to the **Data Synchronization Management** page to view the task status.

Step 2 Manage the migration task.

The migration task contains two phases: full migration and incremental migration. You can manage them in different phases.

- Full synchronization

  Viewing the synchronization progress: Click the target full synchronization task, and on the **Migration Progress** tab, you can see the synchronization progress of the structure, data, indexes, and synchronization objects. When the progress reaches 100%, the synchronization is complete.
- Incremental synchronization
  - Viewing the synchronization delay: After the full synchronization is complete, an incremental synchronization starts. On the **Data Synchronization Management** page, click the target synchronization task. On the displayed page, click **Synchronization Progress** to view the synchronization delay of the incremental synchronization. If the synchronization delay is 0s, the destination database is being

 Copyright © Huawei Technologies Co., Ltd.

synchronized with the source database in real time. You can also view the data consistency on the **Synchronization Comparison** tab.

Figure 4-25 Viewing the synchronization delay

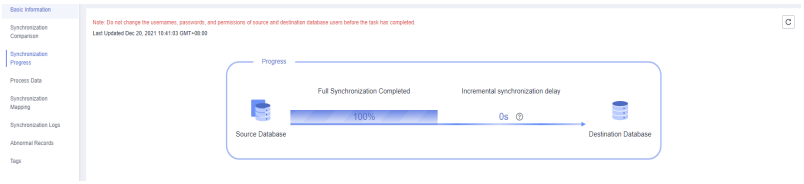

- Viewing the synchronization comparison results: To minimize service downtime, click the name of an incremental synchronization task. On the **Synchronization Comparison** page, create a comparison task.

  For details, see **Comparing Migration Items** in *Data Replication Service User Guide*.

Step 3 Cut over services.

You are advised to start the cutover process during off-peak hours. At least one complete data comparison is performed during off-peak hours. 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, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

1. Interrupt services first. If the workload is not heavy, you may not need to interrupt the services.
2. Run the following statement on the source database and check whether any new sessions execute SQL statements within the next 1 to 5 minutes. If there are no new statements executed, the service has been stopped.

   ```
   select * from pg_stat_activity;
   ```

   NOTE

   The process list queried by the preceding statement includes the connection of the DRS synchronization instance. If no additional session executes SQL statements, the service has been stopped.
3. On the **Synchronization Progress** page, view the synchronization delay. When the delay is displayed as 0s and remains stable for a period, then you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the results of the previous comparison.
   - If there is enough time, compare all objects.
   - If there is not enough time, use the data-level comparison to compare the tables that are frequently used and that contain key business data or inconsistent data.
4. Determine an appropriate time to cut the services over to the destination database. After services are restored and available, the synchronization is complete.

Step 4 Complete the synchronization.

1. Stop the synchronization task. After databases and services are synchronized to the destination database, to prevent operations on the source database

 Copyright © Huawei Technologies Co., Ltd.

from being synchronized to the destination database to overwrite data, you can stop the synchronization task. This operation only deletes the synchronization instance, and the synchronization task is still displayed in the task list. You can view or delete the task. You will not be charged for the synchronization task after the task is stopped.

2. Delete the synchronization task. After the synchronization task is complete, you can delete it. After the synchronization task is deleted, it will not be displayed in the task list.

----End

4.4 From RDS MySQL to Kafka

4.4.1 Overview

Description

In this section, we will create a DRS synchronization task to synchronize the incremental data from the source RDS MySQL database to the destination Kafka database. This section also describes:

- How to create an RDS MySQL instance on Huawei Cloud.
- How to create DMS for Kafka on Huawei Cloud.
- How to create a DRS synchronization task.

Prerequisites

- You have registered with Huawei Cloud.
- Your account balance is greater than or equal to $0 USD.

Service List

- Virtual Private Cloud (VPC)
- Relational Database Service (RDS)
- DMS for Kafka
- Data Replication Service (DRS)
- Data Admin Service (DAS)

Deployment Architecture

In this example, the source is an RDS MySQL instance, and the destination is a DMS for Kafka in the same region. Incremental data of the source database is synchronized to the destination database in a VPC. For details about the deployment architecture, see **Figure 4-26**.

 Copyright © Huawei Technologies Co., Ltd.

Figure 4-26 VPC network

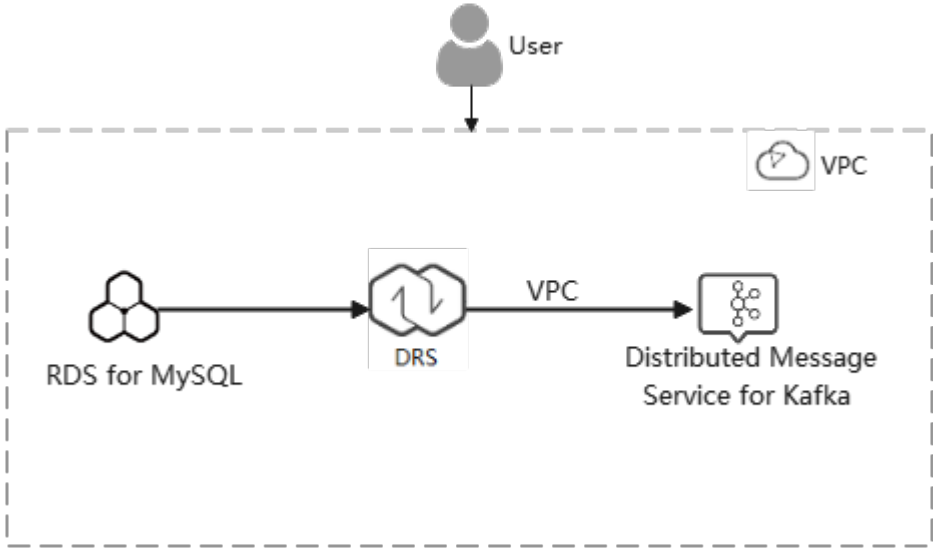

Before You Start

- The resource planning in this best practice is for demonstration only. Adjust it as needed.
- The test data is for reference only. For more information about DRS, click **here**.

4.4.2 Resource and Cost Planning

Table 4-20 Resource planning

Category	Subcategory	Plan	Description
VPC	VPC name	vpc-DRStest	Specify a name that is easy to identify.
	Region	EU-Dublin	To achieve lower network latency, select the region nearest to you.
	AZ	AZ 1	-
	Subnet	10.0.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-drs01	Specify a name that is easy to identify.
RDS (source database)	RDS instance name	rds-mysql	Specify a name that is easy to identify.
	DB engine version	MySQL 5.7	-

Copyright © Huawei Technologies Co., Ltd.

Category	Subcategory	Plan	Description
	Instance type	Single	A single instance is used in this example. To improve service reliability, select a primary/standby instance.
	Storage type	SSD	-
	AZ	AZ 3	A single instance is used in this example. To improve service reliability, create a primary and standby instance and then locate them in two different AZs.
	Specifications	General-purpose 4 vCPUs | 8 GB	-
Kafka (destination database)	Kafka instance name	kafka-drs	Specify a name that is easy to identify.
	Version	2.3.0	-
	AZ	AZ 3	You can select one, three, or more AZs. You are advised to create the instance across different AZs to improve service reliability.
	Specifications	c6.2u4g.cluster	-
	Brokers	3	-
	Storage space	High I/O, 200 GB	The storage space is used to store messages (including replicas). Kafka uses three replicas by default. In addition to storing messages, some space needs to be reserved for storing logs and metadata.
DRS synchronization task	Synchronization task name	DRS-MySQLToKafka	Custom
	Source DB engine	MySQL	In this example, the source is an RDS MySQL instance on Huawei Cloud.
	Destination DB engine	Kafka	In this example, the destination database is Kafka.
	Network Type	VPC	In the practice, select the VPC network.

Copyright © Huawei Technologies Co., Ltd.

NOTE

For details about the prices of the preceding resources, see **Pricing Details**. You can use the **price calculator** provided by Huawei Cloud to quickly calculate the reference price based on your desired specifications.

4.4.3 Operation Process

Figure 4-27 shows the process of creating an RDS MySQL instance and synchronizing the incremental data from an RDS MySQL instance to Kafka.

Figure 4-27 Flowchart

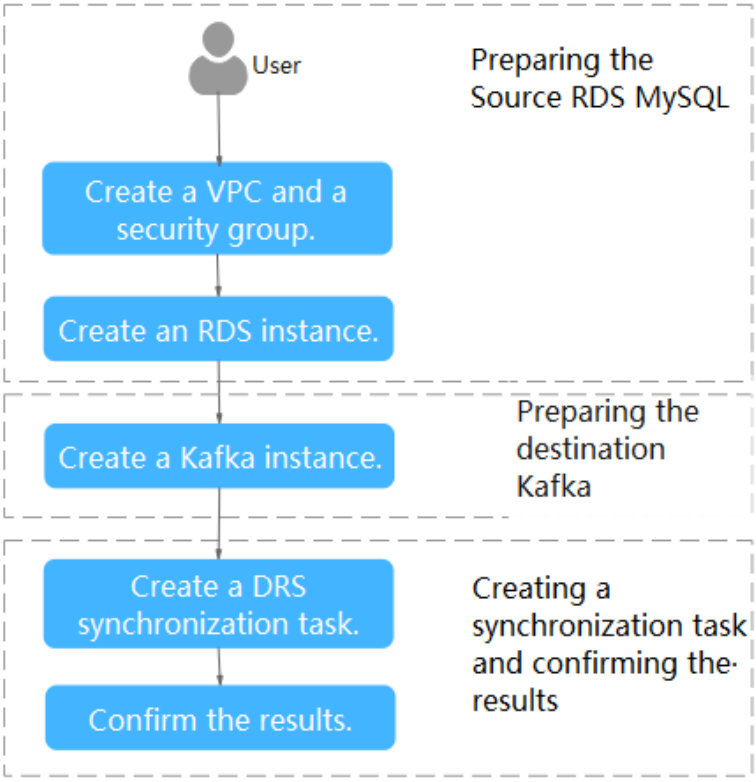

4.4.4 Creating a VPC and Security Group

Create a VPC and security group for an RDS for MySQL and Kafka instances.

Creating a VPC

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select EU-Dublin.

Step 3 Under the service list, choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

 Copyright © Huawei Technologies Co., Ltd.

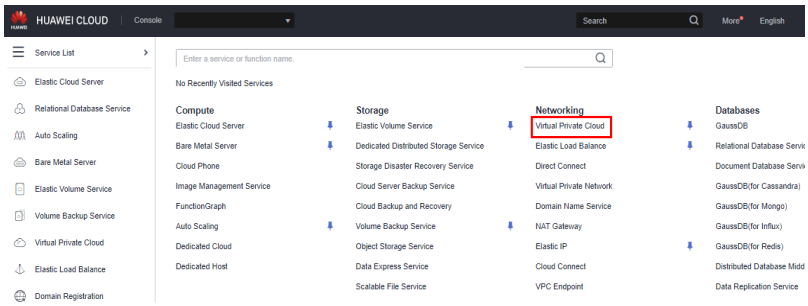

Step 4 Click **Create VPC**.

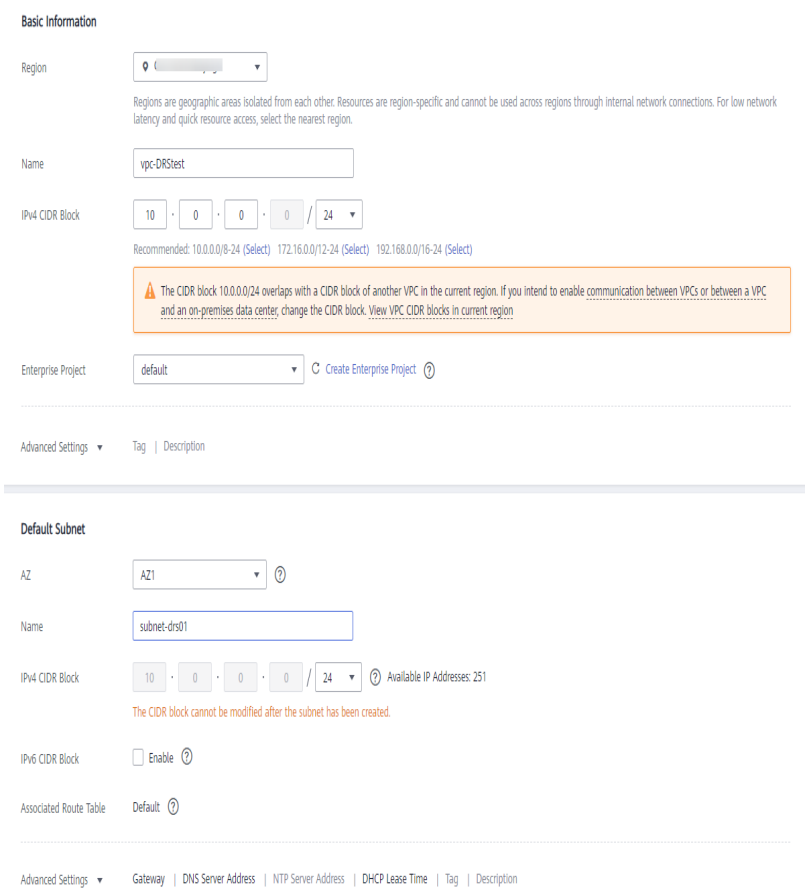

Step 5 Configure parameters as needed and click **Create Now**.

Step 6 Return to the VPC list and check whether the VPC is created.

If the VPC status becomes available, the VPC has been created.

----End

Creating a Security Group

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select EU-Dublin.

 Copyright © Huawei Technologies Co., Ltd.

Step 3 Under the service list, choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

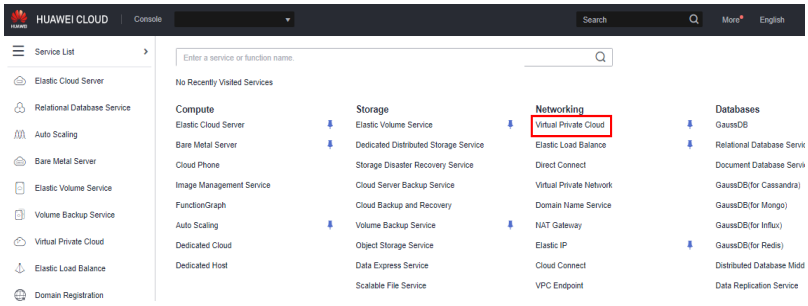

Step 4 In the navigation pane, choose **Access Control** > **Security Groups**.

Step 5 Click **Create Security Group**.

Step 6 Configure parameters as needed.

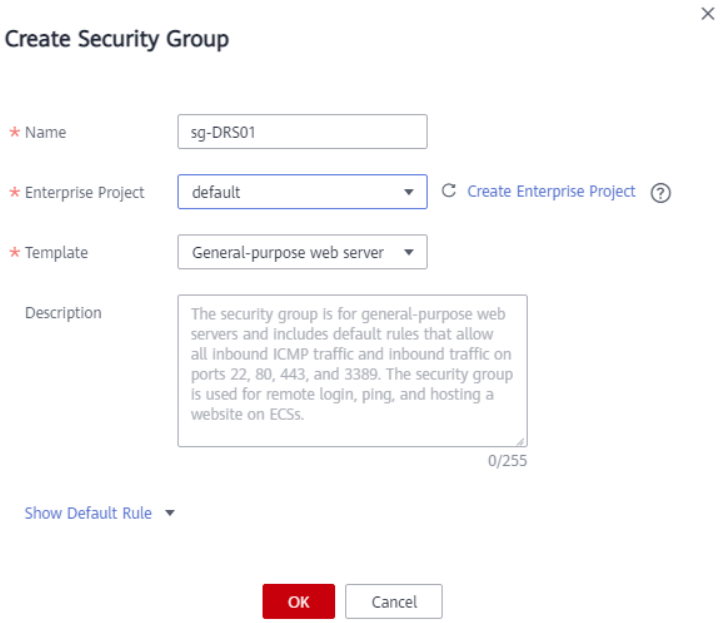

Step 7 Click **OK**.

Step 8 Return to the security group list and click the security group name (**sg-DRS01** in this example).

Step 9 Click the **Inbound Rules** tab, and then click **Add Rule**.

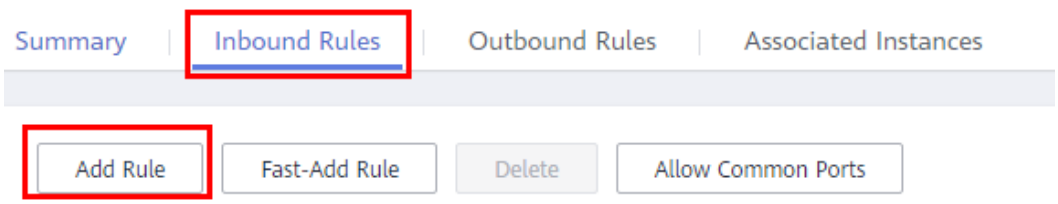

 Copyright © Huawei Technologies Co., Ltd.

Step 10 Configure an inbound rule to allow access from database port **3306**.

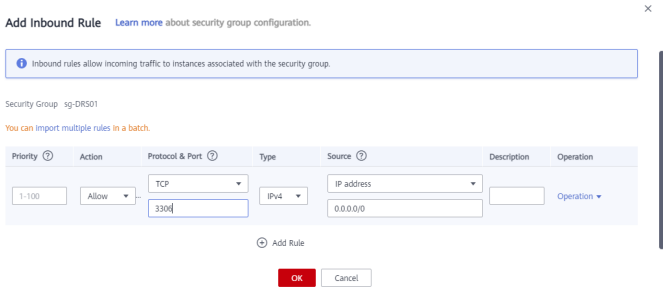

----End

4.4.5 Preparing for Source RDS MySQL

4.4.5.1 Creating an RDS MySQL Instance

Create an RDS MySQL instance, and select the VPC and security group you configured for the instance.

Step 1 Log in to the **management console**.

Step 2 Click  in the upper left corner of the management console and select EU-Dublin.

Step 3 Under the service list, choose **Databases** > **Relational Database Service**.

Step 4 Click **Buy DB Instance**.

Step 5 Configure the instance name and basic information.

Step 6 Configure instance specifications.

 Copyright © Huawei Technologies Co., Ltd.

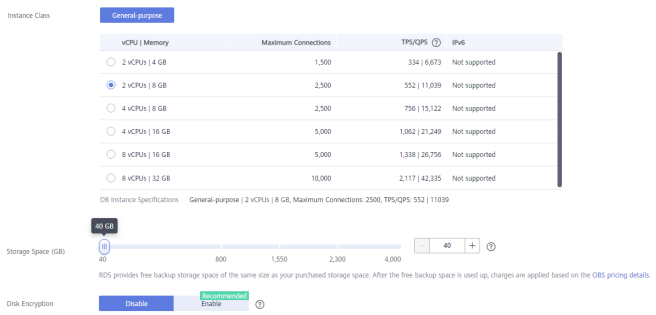

Step 7 Select a VPC and security group for the instance and configure the database port.

The VPC and security group have been created in **Creating a VPC and Security Group**.

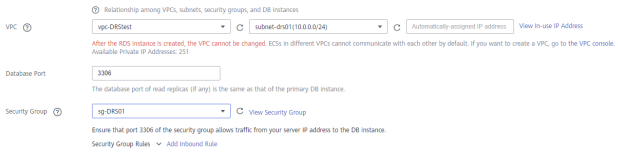

Step 8 Configure the instance password.

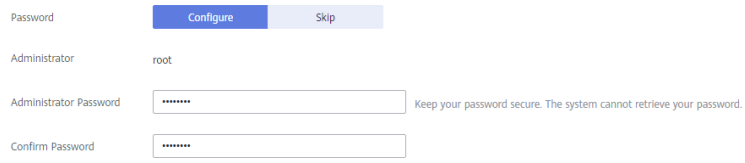

Step 9 Click **Next**.

Step 10 Return to the instance list.

If the instance status becomes available, the instance has been created.

----End

4.4.5.2 Generating Test Data

Step 1 Log in to the **management console**.

Step 2 Click  in the upper left corner of the management console and select EU-Dublin.

Step 3 Under the service list, choose **Databases** > **Relational Database Service**.

Step 4 Select an RDS instance and choose **More** > **Log In**.

Step 5 In the displayed dialog box, enter the password and click **Test Connection**.

Step 6 After the connection is successful, click **Log In**.

Step 7 Enter the instance password for logging in to the RDS instance.

Step 8 Click **Create Database** to create the **db_test** database.

 Copyright © Huawei Technologies Co., Ltd.

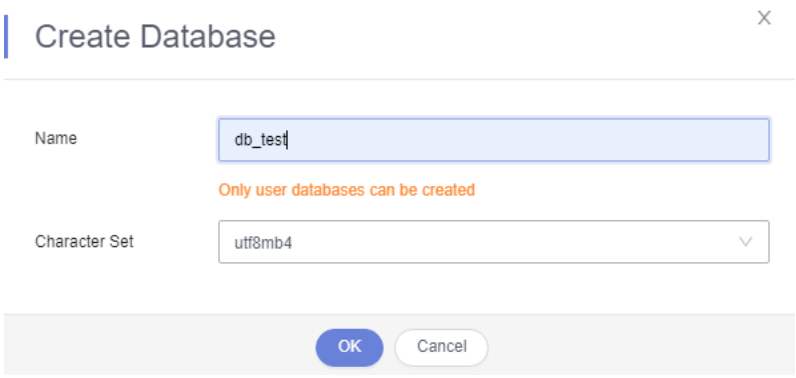

Step 9 Run the following statement in the db_test database to create the corresponding table **table3_**:

```
CREATE TABLE `db_test`.`table3_` (
   `Column1` INT(11) UNSIGNED NOT NULL,
   `Column2` TIME NULL,
   `Column3` CHAR NULL,
   PRIMARY KEY (`Column1`)
)   ENGINE = InnoDB
   DEFAULT CHARACTER SET = utf8mb4
   COLLATE = utf8mb4_general_ci;
```

----End

4.4.6 Preparing for Destination Kafka

4.4.6.1 Creating a Kafka Instance

This section describes how to create a Kafka instance.

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select EU-Dublin.

Step 3 In the Service List, choose **Application** > **Distributed Message Service for Kafka** to open the Kafka console.

 Copyright © Huawei Technologies Co., Ltd.

Step 4 Click **Buy Kafka Instance**.

Step 5 Select the instance region and AZ.

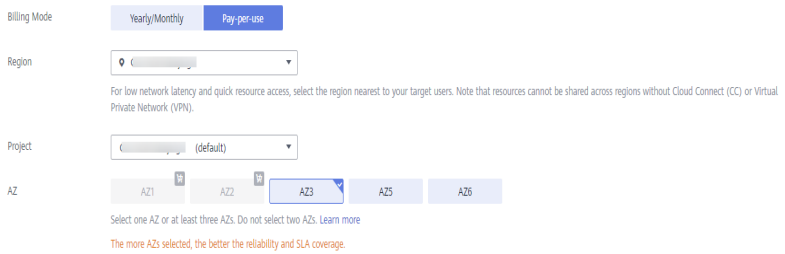

Step 6 Configure the instance name and specifications.

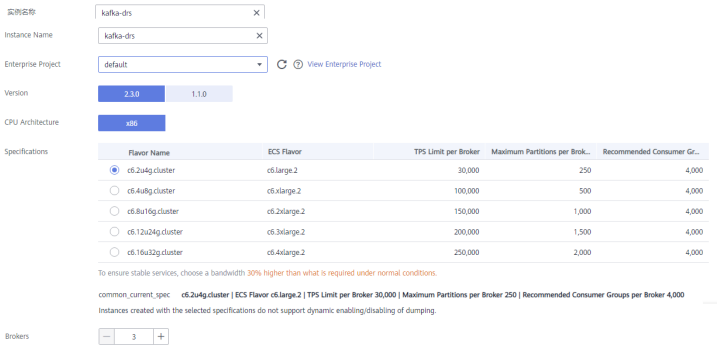

Step 7 Select the storage space and capacity threshold policy.

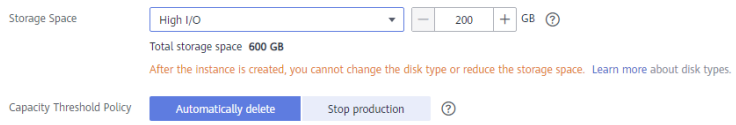

Step 8 Select the VPC and security group.

The VPC and security group have been created in **Creating a VPC and Security Group**.

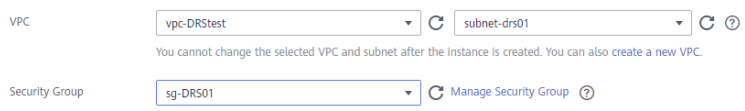

Step 9 Configure the instance password.

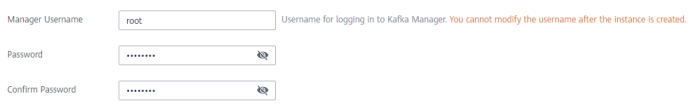

Step 10 Click **Next**.

Step 11 Return to the instance list.

If the status of the Kafka instance is **Running**, the instance has been created.

----End

 Copyright © Huawei Technologies Co., Ltd.

4.4.6.2 Creating a Topic

Step 1 Click a Kafka instance.

Step 2 Click the **Topics** tab, and click **Create Topic**.

Step 3 In the dialog box that is displayed, enter a topic name, specify other parameters, and click **OK**.

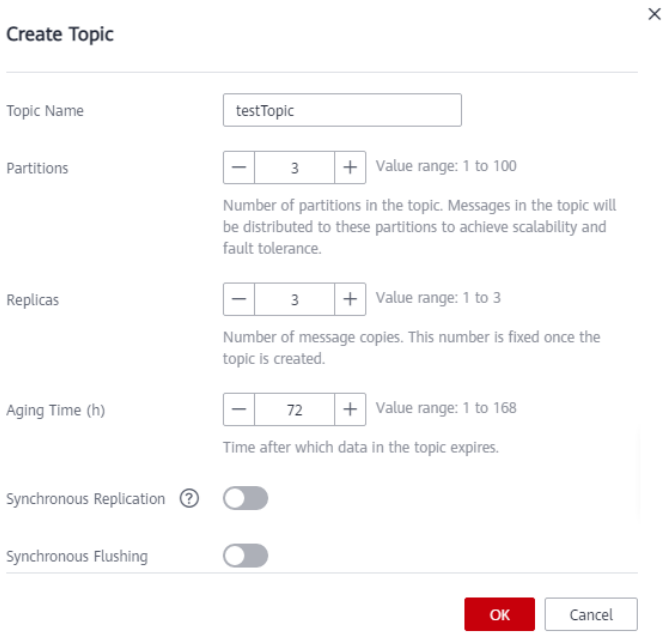

----End

4.4.7 Creating a DRS Synchronization Task

This section describes how to create a DRS instance and synchronize incremental data from RDS MySQL to Kafka.

Pre-Check

Before creating a task, check the synchronization conditions.

In this example, data is synchronized from RDS MySQL to Kafka. For details, see **Precautions**.

Procedure

This section describes how to perform incremental synchronization from RDS MySQL to Kafka.

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select EU-Dublin.

Step 3 Under the service list, choose **Databases** > **Data Replication Service**.

 Copyright © Huawei Technologies Co., Ltd.

Step 4 Choose **Data Synchronization Management** and click **Create Synchronization Task**.

Step 5 Configure synchronization task parameters.

1. Task name

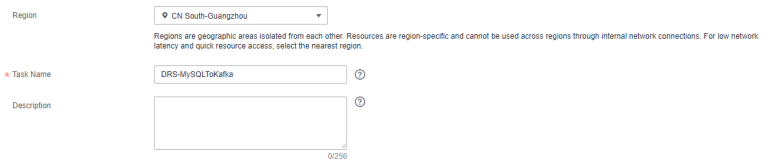

2. Select the source database, destination database, and network information.

   Select the RDS instance created in **Preparing for Source RDS MySQL** as the destination database.

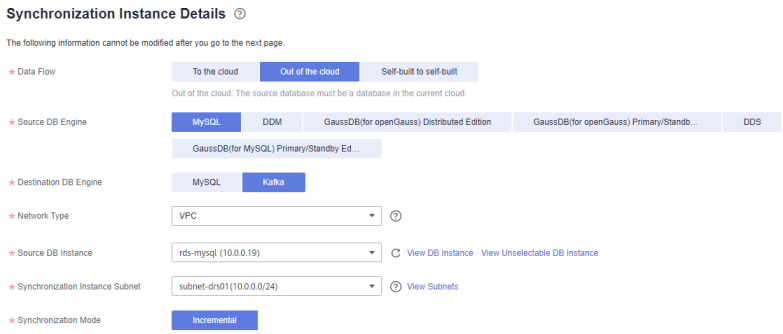

3. Set **Enterprise Project** to **default**.

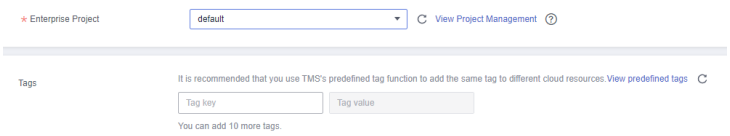

Step 6 Click **Next**.

The synchronization instance is being created. It takes about 5 to 10 minutes.

Step 7 Configure source and destination database information.

1. Configure source database information.
2. Click **Test Connection**.

   If a successful test message is returned, login to the destination is successful.

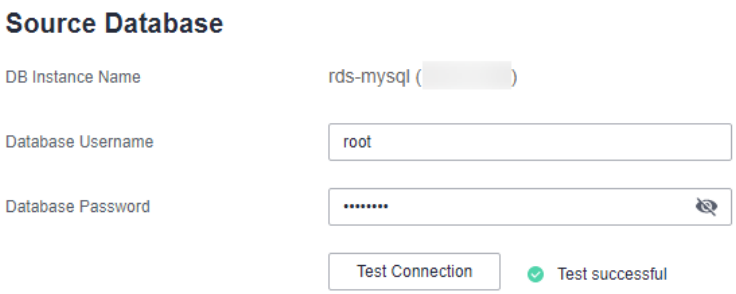

3. Select the VPC and subnet where the destination database is located, and enter the Kafka IP address and port number.

 Copyright © Huawei Technologies Co., Ltd.

4. Click **Test Connection**.

   If a successful test message is returned, login to the destination is successful.

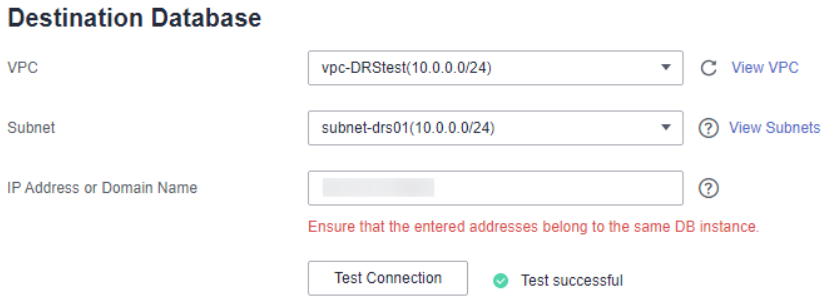

Step 8 Click **Next**.

Step 9 Select the synchronization information, policy, message format, and object, and the format of the message sent to the Kafka.

The following table lists the settings.

Table 4-21 Sync Settings

Type	Setting
Topic Synchronization Policy	Deliver the content to a topic named **testTopic**.
Synchronize Topic To	Partitions are differentiated by the hash values of the database and table names
Data Format in Kafka	You can select the JSON format. For details, see **Kafka Message Format**.
Synchronization Object	Select **table3_** under **db_test**.

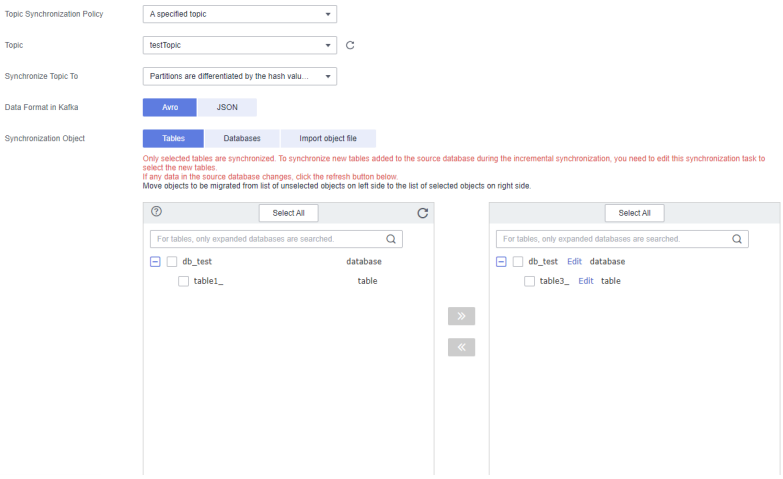

Step 10 Click **Next**.

 Copyright © Huawei Technologies Co., Ltd.

Step 11 Select a data processing mode. Data synchronization from RDS MySQL to Kafka supports only column-based processing, which provides column-level query and filtering capabilities.

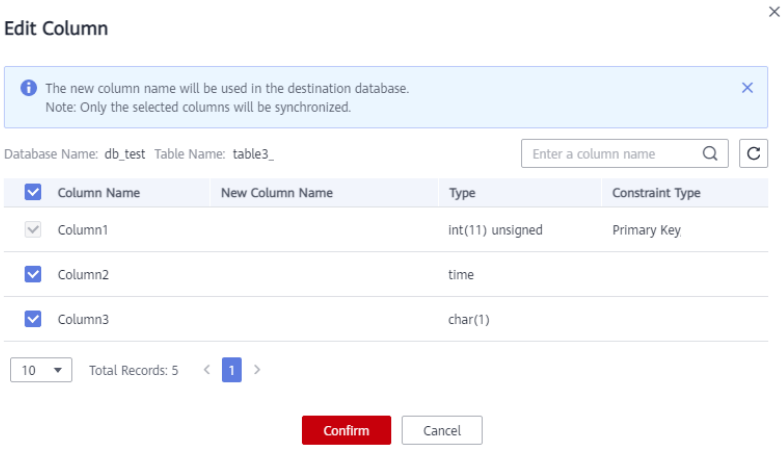

Step 12 Click **Next** and wait for the check results.

Step 13 If the check is complete and the check success rate is 100%, click **Next**.

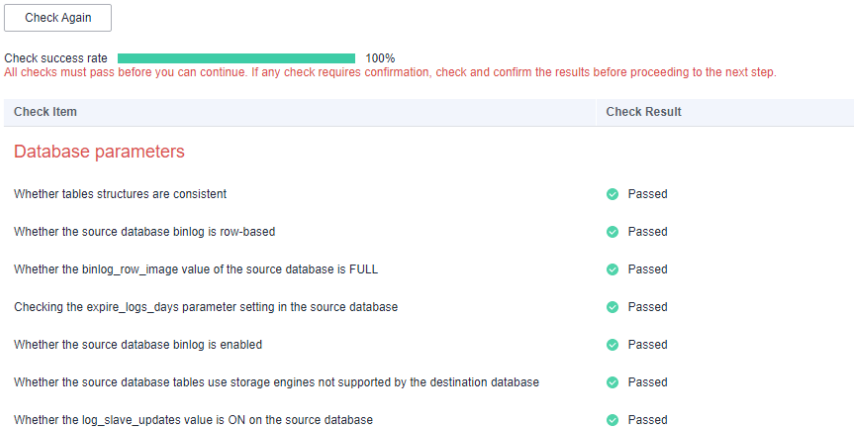

Step 14 After confirming that the synchronization task information is correct, click **Next**.

Return to the **Data Synchronization Management** page and check the synchronization task status.

It takes several minutes to complete.

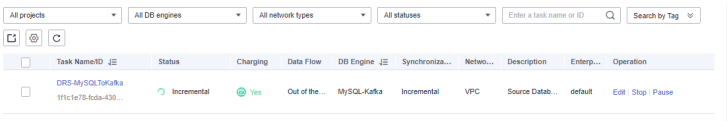

If the status changes to **Incremental synchronization**, the synchronization task has been started.

 Copyright © Huawei Technologies Co., Ltd.

NOTE

- RDS for MySQL supports only incremental synchronization to Kafka. After the task is started, it is in the incremental synchronization state.
- If the created task is a full synchronization task, full synchronization is performed after the task is started. After the synchronization is complete, the task automatically stops.
- If the created task is a full+incremental synchronization task, the task enters the full synchronization state after being started. After the full synchronization is complete, the task enters the incremental synchronization state.
- Incremental synchronization continuously synchronizes incremental data and does not stop automatically.

----End

4.4.8 Confirming the Results

In this practice, DRS continuously synchronizes the incremental data generated in the source database to the destination database until you stop the task. The following describes how to verify the synchronization results by inserting data to the source RDS MySQL database and viewing the data received by Kafka.

Procedure

Step 1 Log in to the **management console**.

Step 2 Click in the upper left corner of the management console and select EU-Dublin.

Step 3 Under the service list, choose **Databases** > **Relational Database Service**.

Step 4 Locate the required RDS instance and choose **More** > **Log In**.

Step 5 In the displayed dialog box, enter the password and click **Test Connection**.

Step 6 After the connection is successful, click **Log In**.

Step 7 Enter the instance password for logging in to the RDS instance.

Step 8 Run the following statement to insert data to the **db_test.table3_** table.

```
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(4,'00:00:44','ddd');
```

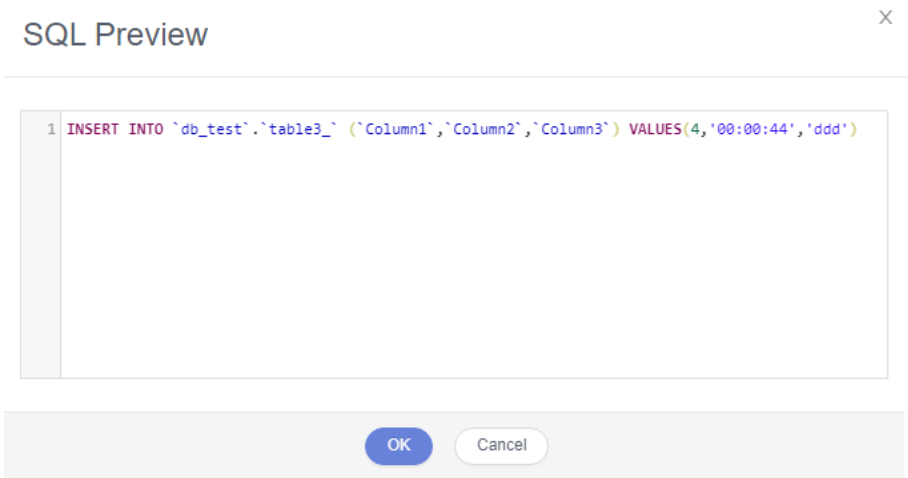

 Copyright © Huawei Technologies Co., Ltd.

Step 9 In the Service List, choose **Application** > **Distributed Message Service for Kafka** to open the Kafka console.

Step 10 Click a Kafka instance.

Step 11 Click the **Message Query** tab, select the Kafka topic, view the received data in JSON format.

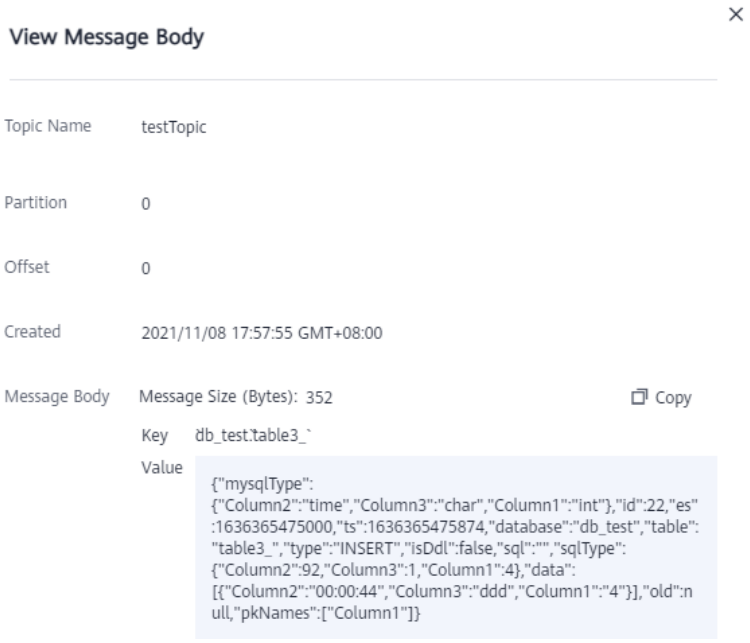

Step 12 Stop the synchronization task.

If all data has been synchronized to the destination database, you can stop the current task.

1. Locate the task and click **Stop** in the **Operation** column.

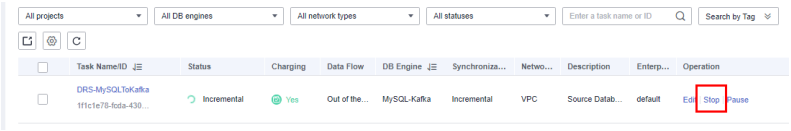

2. In the display box, click **Yes**.

 Copyright © Huawei Technologies Co., Ltd.

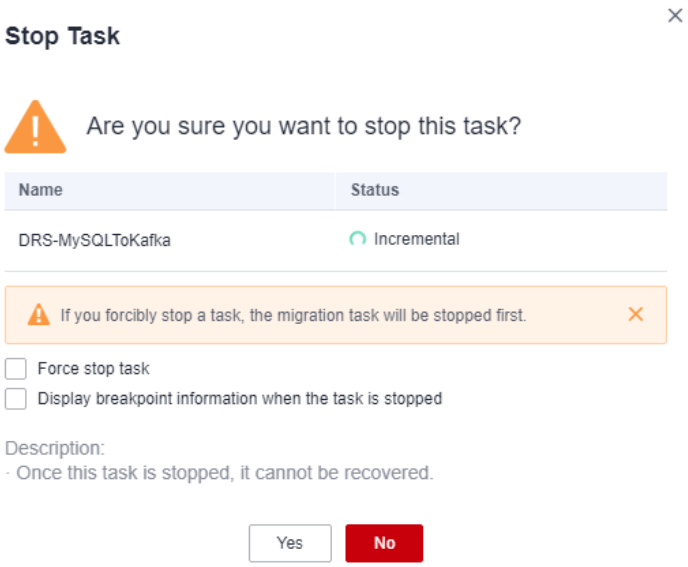

----End

 Copyright © Huawei Technologies Co., Ltd.

A Change History

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

 Copyright © Huawei Technologies Co., Ltd.