Data Replication Service

Best Practices

Issue 17

Date 2025-09-05





Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.

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

Trademarks and Permissions

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

Notice

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

Contents

1 Best Practices Summary	1
2 Real-Time Migration	4
2.1 Migrating MySQL Databases from Other Clouds to RDS for MySQL	
2.1.1 Overview	4
2.1.2 Resource Planning	5
2.1.3 Operation Process	6
2.1.4 Creating a VPC and Security Group	6
2.1.5 Creating an RDS for MySQL Instance	3
2.1.6 Configuring a MySQL Instance on Another Cloud	
2.1.7 Cloud Migration	12
2.1.7.1 Creating a DRS Migration Task	12
2.1.7.2 Checking Migration Results	14
2.2 From Other Cloud MySQL to TaurusDB	15
2.2.1 Overview	15
2.2.2 Resource Planning	16
2.2.3 Operation Process	18
2.2.4 Creating a VPC and Security Group	18
2.2.5 Creating a TaurusDB Instance	19
2.2.6 Configuring a MySQL Instance on Other Clouds	21
2.2.7 Creating a DRS Migration Task	22
2.2.8 Checking Migration Results	24
2.3 From Other Cloud MongoDB to DDS	25
2.4 From ECS-hosted MySQL to RDS for MySQL	40
2.4.1 Overview	40
2.4.2 Resource Planning	42
2.4.3 Operation Process	42
2.4.4 Cloud Migration	43
2.4.4.1 Creating an RDS for MySQL Instance	43
2.4.4.2 Creating a Migration Task	45
2.4.4.3 Confirming Migration Results	47
2.5 From ECS-hosted MySQL to TaurusDB	48
2.5.1 Overview	48
2.5.2 Prerequisites	51

2.5.3 Resource Planning	51
2.5.4 Operation Process	52
2.5.5 Procedure	53
2.5.5.1 ECS-hosted MySQL Server	53
2.5.5.1.1 Creating a VPC and Security Group	53
2.5.5.1.2 Creating an ECS (MySQL Server)	55
2.5.5.1.3 Installing a MySQL Database (Community Edition)	58
2.5.5.1.4 Creating an ECS and Installing a MySQL Client on It	60
2.5.5.1.5 Testing the ECS-hosted MySQL Performance	60
2.5.5.2 Cloud Migration	62
2.5.5.2.1 Creating a TaurusDB Instance	62
2.5.5.2.2 Creating a DRS Migration Task	63
2.5.5.2.3 Checking the Migration Results	65
2.6 From ECS-hosted MongoDB to DDS	66
2.7 From On-Premises MySQL to RDS for MySQL	81
2.8 From On-Premises MongoDB to DDS	96
2.9 From RDS for MySQL to DDM	111
2.9.1 Overview	111
2.9.2 Resource Planning	113
2.9.3 Operation Process	116
2.9.4 Preparing for the Source RDS for MySQL Instance	117
2.9.4.1 Creating a VPC and Security Group	117
2.9.4.2 Creating an RDS for MySQL Instance	120
2.9.5 Preparing for the Destination DDM Instance	123
2.9.5.1 Creating a VPC and Security Group	123
2.9.5.2 Creating a DDM Instance	125
2.9.5.3 Creating an RDS for MySQL Instance	126
2.9.5.4 Creating a Schema and Associating It with the RDS for MySQL Instance	128
2.9.5.5 Creating a DDM Account	129
2.9.5.6 Creating Table Structures in the Destination Database	130
2.9.6 Creating and Configuring a VPN	131
2.9.6.1 Creating a VPN for the Source Database	131
2.9.6.2 Creating a VPN for the Destination Database	133
2.9.6.3 Modifying the VPN Configuration for the Source Database	134
2.9.7 Creating a DRS Migration Task	135
2.9.8 Confirming Migration Results	137
2.10 From MySQL Schema and Logic Table to DDM	139
2.10.1 Overview	139
2.10.2 Resource Planning	140
2.10.3 Process	
2.10.4 Procedure	143
2.10.4.1 Constructing Data Before the Migration	143

2.10.4.2 Preparing for the Destination DDM Instance	145
2.10.4.2.1 Creating a VPC and Security Group	145
2.10.4.2.2 Creating a DDM Instance	
2.10.4.2.3 Creating an RDS for MySQL Instance	149
2.10.4.2.4 Creating a DDM Account	151
2.10.4.2.5 Creating a Schema and Associating It with the RDS for MySQL Instance	151
2.10.4.2.6 Creating Table Structures in the Destination Database	152
2.10.4.3 Migrating Databases	154
2.10.4.4 Verifying Data After Migration	158
3 Backup Migration	161
3.1 Migrating Microsoft SQL Server Backup Data to RDS for SQL Server	161
3.1.1 Overview	161
3.1.2 Migration Preparations	162
3.1.3 Exporting Backup Files	164
3.1.4 Uploading Backup Files	170
3.1.5 Scenario 1: Full Backup Migration	171
3.1.6 Scenario 2: Full+Incremental Backup Migration	175
3.1.7 Manual Configuration	188
4 Real-Time Synchronization	196
4.1 From Other Cloud PostgreSQL to RDS for PostgreSQL	
4.2 From ECS-hosted PostgreSQL to RDS for PostgreSQL	208
4.3 From On-Premises PostgreSQL to RDS for PostgreSQL	222
4.4 From On-premises MySQL to GaussDB Distributed	234
4.5 From On-premises Oracle to GaussDB Distributed	255
4.5.1 Purpose	255
4.5.2 Resource Planning	257
4.5.3 Operation Process	259
4.5.4 Creating a VPC and Security Group	260
4.5.5 Creating a GaussDB Instance	263
4.5.6 Generating Test Data	265
4.5.7 Synchronizing Databases	277
4.5.8 Verifying Data	284
4.6 From On-premises Oracle to GaussDB Centralized	287
4.6.1 Purpose	287
4.6.2 Resource Planning	289
4.6.3 Operation Process	292
4.6.4 Preparations	293
4.6.5 Creating a VPC and Security Group	295
4.6.6 Creating a GaussDB Instance	298
4.6.7 Creating a UGO Task	300
4.6.7.1 Creating a Database Evaluation Project	300
4.6.7.2 Creating an Object Migration Project	303

4.6.9 Creating a DDC Took	204
4.6.8 Creating a DRS Task	
4.6.9.1 UGO Schema Migration Phase 1: Tables, Primary Keys, Unique Keys/Indexes	
4.6.9.2 DRS Synchronization Phase 1: Full Synchronization	
4.6.9.3 UGO Schema Migration Phase 2: Common Indexes	
4.6.9.4 DRS Synchronization Phase 2: Incremental Synchronization	
4.6.9.5 UGO Schema Migration Phase 3: Triggers, Events, Tasks, Foreign Keys, and Sequences	
4.7 From On-Premises Oracle to DDM	
4.7.2 Resource Planning	
4.7.4 Description Process	
4.7.4 Preparing for the Source Oracle Database	
4.7.5 Preparing for the Destination DDM Instance	
4.7.5.1 Creating a VPC and Security Group	
4.7.5.2 Creating a DDM Instance	
4.7.5.3 Creating an RDS for MySQL Instance	
4.7.5.4 Creating a Schema and Associating It with the RDS for MySQL Instance	
4.7.5.5 Creating a DDM Account	326
4.7.5.6 Creating Table Structures in the Destination Database	327
4.7.6 Creating a DRS Synchronization Task	327
4.7.7 Confirming Synchronization Results	330
4.8 From RDS for MySQL to Kafka	332
4.8.1 Overview	332
4.8.2 Resource and Cost Planning	333
4.8.3 Operation Process	335
4.8.4 Creating a VPC and Security Group	335
4.8.5 Preparing for Source RDS for MySQL	338
4.8.5.1 Creating an RDS for MySQL Instance	338
4.8.5.2 Generating Test Data	339
4.8.6 Preparing for Destination Kafka	340
4.8.6.1 Creating a Kafka Instance	341
4.8.6.2 Creating a Topic	
4.8.7 Creating a DRS Synchronization Task	
4.8.8 Confirming the Results	
4.9 From GaussDB Distributed to Kafka	
4.10 From GaussDB Centralized to Kafka	
5 Real-Time Disaster Recovery	
5.1 Configuring Remote Single-Active DR for an RDS for MySQL Instance Using DRS	
5.1.1 Overview	
5.1.2 Resource Planning	
5.1.3 Operation Process	
5.1.4 Configuring an RDS for MySQL Instance in the Production Center	

6 Security Best Practices	401
5.1.6.3 Performing a Primary/Standby Switchover	399
5.1.6.2 Configuring the Disaster Recovery Task	397
5.1.6.1 Creating a DRS Disaster Recovery Task	
5.1.6 Configuring Remote Disaster Recovery	
5.1.5.2 Creating an RDS for MySQL Instance	
5.1.5.1 Creating a VPC and Security Group	392
5.1.5 Configuring an RDS for MySQL Instance in the DR Center	392
5.1.4.3 Creating an RDS for MySQL Instance	389
5.1.4.2 Creating an EIP	389
5.1.4.1 Creating a VPC and Security Group	

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 Migratio	MySQL databases on other clouds	Huawei Cloud RDS for MySQL instances	Migrating MySQL Databases from Other Clouds to RDS for MySQL
n		Huawei Cloud TaurusDB instances	From Other Cloud MySQL to TaurusDB
	MongoDB databases on other clouds	Huawei Cloud DDS instances	From Other Cloud MongoDB to DDS
	On-premises MySQL databases	Huawei Cloud RDS for MySQL instances	From ECS-hosted MySQL to RDS for MySQL
	MySQL databases on ECSs	Huawei Cloud TaurusDB instances	From ECS-hosted MySQL to TaurusDB
	MongoDB databases on ECSs	Huawei Cloud DDS instances	From ECS-hosted MongoDB to DDS
	On-premises MySQL databases	Huawei Cloud RDS for MySQL instances	From On-Premises MySQL to RDS for MySQL
	On-premises MongoDB databases	Huawei Cloud DDS instances	From On-Premises MongoDB to DDS

Function	Source DB Type	Destination DB	Cases
	RDS for MySQL instances	Distributed Database Middleware (DDM) instances	From RDS for MySQL to DDM
	MySQL schema and logic table	Distributed Database Middleware (DDM) instances	From MySQL Schema and Logic Table to DDM
Backup Migratio n	On-premises Microsoft SQL Server databases	Huawei Cloud RDS for SQL Server instances	Migrating Microsoft SQL Server Backup Data to RDS for SQL Server
Real- Time Synchron	PostgreSQL databases on other clouds	Huawei Cloud RDS for PostgreSQL instances	From Other Cloud PostgreSQL to RDS for PostgreSQL
ization	PostgreSQL databases on ECSs		From ECS-hosted PostgreSQL to RDS for PostgreSQL
	On-premises PostgreSQL databases		From On-Premises PostgreSQL to RDS for PostgreSQL
	On-premises MySQL databases	Huawei Cloud distributed GaussDB instances	From On-premises MySQL to GaussDB Distributed
	On-premises Oracle databases	Huawei Cloud distributed GaussDB instances	From On-premises Oracle to GaussDB Distributed
		Huawei Cloud centralized GaussDB instances	From On-premises Oracle to GaussDB Centralized
		Distributed Database Middleware (DDM) instances	From On-Premises Oracle to DDM
	Huawei Cloud RDS for MySQL instances	DMS for Kafka	From RDS for MySQL to Kafka

Function	Source DB Type	Destination DB	Cases
	GaussDB Distributed	DMS for Kafka	From GaussDB Distributed to Kafka
	Huawei Cloud centralized GaussDB instances	DMS for Kafka	From GaussDB Centralized to Kafka
Real- Time Disaster Recovery	Huawei Cloud RDS for MySQL instances	Huawei Cloud RDS for MySQL instances	Configuring Remote Single- Active DR for an RDS for MySQL Instance Using DRS

2 Real-Time Migration

2.1 Migrating MySQL Databases from Other Clouds to RDS for MySQL

2.1.1 Overview

Scenarios

This best practice includes the following tasks:

- Create an RDS for MySQL instance.
- Migrate data from a MySQL database on other clouds to RDS for MySQL.

Prerequisites

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

Service List

- Virtual Private Cloud (VPC)
- RDS
- 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 MySQL.

2.1.2 Resource Planning


Table 2-1 Resource planning

Categor y	Subcatego ry	Planned Value	Description
VPC	VPC name	vpc-src-172	Specify a name that is easy to identify.
	Region	Test region	To achieve lower network latency, select the region nearest to you.
	AZ	AZ3	-
	Subnet	172.16.0.0/16	Select a subnet with sufficient network resources.
	Subnet name	subnet-src-172	Specify a name that is easy to identify.
MySQL on	Database version	MySQL 5.7	-
another cloud	IP address	10.154.217.42	Enter an IP address.
	Port	3306	-
RDS for MySQL	Instance name	rds-mysql	Specify a name that is easy to identify.
instance	DB engine version	MySQL 5.7	-
	Instance type	Single	A single instance is used in this example. To improve service reliability, select a primary/ standby instance.
	Storage type	Cloud SSD	-
	AZ	AZ1	AZ1 is selected in this example. To improve service reliability, select the primary/standby instance type and deploy the primary and standby instances in different AZs.
	Instance class	General-purpose 2 vCPUs 8 GB	-
DRS migratio n task	Task name	DRS-mysql	Specify a name that is easy to identify.

Categor y	Subcatego ry	Planned Value	Description
	Source DB engine	MySQL	-
	Destinatio n DB engine	MySQL	-
	Network type	Public network	Public network is used in this example.

2.1.3 Operation Process

Figure 2-1 Flowchart



2.1.4 Creating a VPC and Security Group


Create a VPC and security group for an RDS for MySQL instance

Creating a VPC

Step 1 Go to the **Create VPC** page.

Step 2 Configure the basic information, subnet, and IP address.

Figure 2-2 Creating a VPC



- Step 3 Click Create Now.
- **Step 4** 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 CN-Hong Kong.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**.
- **Step 4** In the navigation pane, choose **Access Control** > **Security Groups**.
- Step 5 Click Create Security Group.
- **Step 6** Configure parameters as needed.

Figure 2-3 Creating a security group



- Step 7 Click Create Now.
- **Step 8** Return to the security group list and click the security group name.
- **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**.

Figure 2-4 Inbound rules




----End

2.1.5 Creating an RDS for MySQL Instance

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


- **Step 1** Go to the **Buy DB Instance** page.
- **Step 2** Configure basic information for the instance. Select **CN-Hong Kong** for **Region**.

Figure 2-5 Basic information



Step 3 Select an instance class.

Figure 2-6 Instance class



Step 4 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**.




Figure 2-7 Network configurations

Step 5 Complete advanced settings.




Figure 2-8 Additional options

Step 6 Click Next.

- **Step 7** Confirm the settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.
- **Step 8** Return to the instance list. If the instance status becomes available, the instance has been created.
 - ----End

2.1.6 Configuring a MySQL Instance on Another Cloud

Prerequisites

- You have purchased a MySQL instance from another cloud vendor platform.
- Your account has the migration permissions listed in Permission Requirements.

Permission Requirements

Table 2-2 lists the permissions required for migrating data from a MySQL instance on another cloud to RDS for MySQL using DRS. For details about the permissions, see **Which MySQL Permissions Are Required for DRS?**

Table 2-2 Migration permissions

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

Network Configuration

You need to enable public accessibility for the source database.

Whitelist Settings

The EIP of the DRS replication instance must be on the whitelist of the source database for the connectivity between the DRS replication instance and the source database. To obtain the EIP of the DRS replication instance, see **Step 3** in **Creating a DRS Migration Task**. This method of configuring a whitelist varies depending on the cloud database vendors. For details, see their official documents.

2.1.7 Cloud Migration

2.1.7.1 Creating a DRS Migration Task

Creating a Migration Task

- **Step 1** Go to the **Create Migration Task** page.
- **Step 2** Configure parameters as needed.
 - 1. Enter the migration task name. Select the region hosting the destination DB instance for **Region**.

Figure 2-9 Migration task



2. Configure the replication instance information.

Select the RDS instance created in **Creating an RDS for MySQL Instance** as the destination database.

Figure 2-10 Replication instance details


Step 3 Click Create Now.

It takes about 5 to 10 minutes to create a replication instance. After the replication instance is created, you can obtain its EIP.



Step 4 Configure the source and destination database information.

Figure 2-11 Configuring the source and destination databases



- Step 5 Click Next.
- **Step 6** On the **Set Task** page, configure parameters as required.
 - Set Flow Control to No.
 - Set Migration Object to All.
- **Step 7** 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 8** Compare source and destination database parameters.
 - If you do not want to compare the parameters, click **Next** to skip this step.
 - If there are inconsistent common parameter values, click **Save Change** to change the destination database values to match those of the source database.
- **Step 9** Click **Submit** to submit the task.

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 is complete.

----End

2.1.7.2 Checking Migration Results

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

- Use DRS to compare migration objects, users, and data of source and destination databases and obtain the migration results. For details, see Checking the Migration Results on the DRS Console.
- 2. Log in to the destination instance to check whether the databases, tables, and data are migrated. For details, see **Checking the Migration Results on the RDS Console**.

Checking the Migration Results on the DRS Console

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select your region.
- **Step 3** Under the service list, choose **Databases** > **Data Replication Service**.
- **Step 4** Click the DRS instance name.
- **Step 5** Click **Migration Comparison** in the navigation pane. Under the **Object-Level Comparison** tab, click **Compare** to check whether all objects have been migrated to the destination instance.
- **Step 6** Click the **Data-Level Comparison** tab. On the displayed page, click **Create Comparison Task** to check whether the databases and tables of the source and destination instances are the same.

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 RDS Console

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select your region.
- Step 3 Click the service list icon on the left and choose Databases > Relational Database Service.
- **Step 4** Locate the destination instance and click **Log In** in the **Operation** column.
- **Step 5** In the displayed dialog box, enter the password and click **Test Connection**.
- **Step 6** After the connection test is successful, click **Log In**.
- **Step 7** Check whether the databases and tables of the source instance have been migrated.

----End

Performing a Performance Test

After the migration is complete, you can perform a performance test as required.

2.2 From Other Cloud MySQL to TaurusDB

2.2.1 Overview

Description

This section includes the following content:

- Create a TaurusDB instance.
- Migrate data from MySQL on other clouds to TaurusDB.

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 TaurusDB instance. Data is migrated from the source to the destination over a public network. For details about the deployment architecture, see **Figure 2-12**.

Other cloud

Region

VPC

Configure the whitelist.

Database EIP

Database service

HUAWEI CLOUD

Region

Mutually accessible in the same subnet by default same same subnet by default same subnet by default same same subnet by default same subnet same subnet by default same subnet same subn

Figure 2-12 Deployment architecture

Service List

- Virtual Private Cloud (VPC)
- TaurusDB
- 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.

2.2.2 Resource Planning


Table 2-3 Resource planning

Categor y	Subcatego ry	Plan	Description
VPC	VPC name	vpc-DRStest	Specify a name that is easy to identify.
	Region	AP-Singapore	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.

Categor y	Subcatego ry	Plan	Description
Other cloud MySQL	DB engine version	MySQL 5.7	-
	IP address	10.154.217.42	Enter an IP address.
	Port	3306	-
TaurusD B 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 migratio n task	Task name	DRS-test-migrate	Specify a name that is easy to identify.
	Source DB engine	MySQL	-
	Destinatio n DB engine	TaurusDB	-
	Network type	Public network	Public network is used in this example.

2.2.3 Operation Process

Figure 2-13 Flowchart



2.2.4 Creating a VPC and Security Group

Create a VPC and security group for a TaurusDB 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 AP-Singapore.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**.
- 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 AP-Singapore.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**.
- **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.




----End


2.2.5 Creating a TaurusDB Instance

This section describes how to create a Huawei Cloud TaurusDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select region AP-Singapore.
- **Step 3** Under the service list, choose **Databases** > **TaurusDB**.
- **Step 4** On the **Instances** page, click **Buy DB Instance**.
- **Step 5** Configure the instance name and basic information.



Step 6 Configure instance specifications.



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 Configure an enterprise project.



- Step 10 Click Next. If you do not need to modify your settings, click Submit.
- **Step 11** Return to the instance list. If the instance becomes **Available**, the instance has been created.

----End

2.2.6 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 TaurusDB instance, the following permissions are required.

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.

2.2.7 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 TaurusDB instance.


Pre-migration Check

Before creating a migration task, check the migration environment.

This section describes how to migrate data from a MySQL database to TaurusDB. 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 AP-Singapore.
- **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.



2. Configure replication instance details as needed.

Select the TaurusDB instance created in **Creating a TaurusDB Instance** as the destination database.

Figure 2-14 Replication instance information



Step 6 Click Create Now.

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.
 - 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.

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 TaurusDB 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

2.2.8 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**.
- 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 TaurusDB 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 AP-Singapore.
- **Step 3** Under the service list, choose **Databases** > **Data Replication Service**.
- Step 4 Click the DRS instance name.
- **Step 5** Choose **Migration Comparison** and select **Object-Level Comparison** to check whether database objects are missing.
- **Step 6** Click **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 TaurusDB Console

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select region AP-Singapore.
- **Step 3** Under the service list, choose **Databases** > **TaurusDB**.

Step 4 On the **Instances** page, 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

2.3 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.

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.

Resource Planning


Table 2-5 Resource planning

Categor y	Subcatego ry	Planned Value	Description
VPC	VPC name	vpc-dds	Specify a name that is easy to identify.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	AZ	AZ1	-
	Subnet	10.0.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-default	Specify a name that is easy to identify.

Categor y	Subcatego ry	Planned Value	Description
MongoD B databas e on other clouds	Database version	MongoDB 4.4	-
	IP address	192.168.0.1	This is only an example.
	Port	8635	This is only an example.
DDS instance	Instance name	dds-test	Specify a name that is easy to identify.
	Database version	DDS 4.4	-
	AZ type	Single-AZ	This is only an example.
	AZ	AZ1	AZ1 is selected in this example.
	Instance specificatio ns	Enhanced II 4 vCPUs 16 GB	The specifications in the left column are selected in this example. You need to select specifications meeting your workload requirements.
DRS migratio n task	Task name	DRS-test-migrate	Specify a name that is easy to identify.
	Source DB engine	MongoDB	-
	Destinatio n DB engine	DDS	-
	Network type	Public network	Public network is used in this example.

Diagram

Figure 2-15 Migrating MongoDB databases from other clouds






Figure 2-16 Migrating MongoDB databases from other cloud servers

Migration Process

Check the full migration progress.

Check the incremental migration progress.

Migrate services.

End

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*.

Preparations

1. Permissions

Table 2-6 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 2-6 Required 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 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 source shard node user must have the readAnyDatabase permission for the source shard node user must have 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.	

Source database permissions:

The source MongoDB database user must have all the required permissions listed in **Table 2-6**. 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

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 2-18.

Figure 2-18 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 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.

- 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 2-19 Replication instance information



Table 2-7 Task settings

Parameter	Description		
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.		
Project	he project corresponds to the current region and can be hanged.		
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 2-8 Replication instance settings

Parameter	Description
Data Flow	To the cloud

Parameter	Description			
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 nonsystem 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.			
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. 			

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 2-20 Source database information




Table 2-9 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 Huawei Cloud 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 2-21 Destination database information



Table 2-10 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 2-22 Migration object




Table 2-11 Migration object

Paramete r	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. 	

Paramete r	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 \mathbb{C} 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.

Check Again

Check success rate

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

Check Result

Destination database storage space

Whether the destination database has sufficient storage space

Conflict

Whether collections in both the source and destination databases are not capped

Whether the destination database contains a non-empty collection with the same name as that in the source database

Whether the same view names exist in both the source and destination databases

Object dependency

Whether the source database referenced roies pass the check

Passed

Object dependency

Whether the source database referenced accounts pass the check

Passed

Database parameters

Whether both the source and destination databases is sufficient

Whether the maximum number of chunks in the destination database is sufficient

Whether the maximum number of chunks in the destination database is sufficient

Whether the maximum number of collections has been reached in the destination database

Passed

Whether the maximum number of collections has been reached in the destination database

Passed

Figure 2-23 Task Check

Ⅲ 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.

 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 2-24 Task startup settings




Table 2-12 Task startup settings

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


- 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 2-25 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 2-26 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.
- 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()

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.

- 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

2.4 From ECS-hosted MySQL to RDS for MySQL

2.4.1 Overview

Scenarios

This chapter includes the following content:

How to migrate data from self-managed MySQL databases to RDS for MySQL instances

RDS for MySQL Advantages

More Services at Lower Costs

You pay for only RDS instances. There is no hardware or management investment needed.

• Ultimate User Experience

- Fully compatible with MySQL
- Excellent performance for high concurrency
- Support for a great number of connections and quicker response

High Security

- End-to-end database security, including network isolation, access control, transmission encryption, storage encryption, and anti-DDoS
- Highest-level certification by the NIST-CSF, with 108 key security capabilities

High Reliability

Multiple deployment and DR solutions, including data backup, data restoration, dual-host hot standby, remote DR, and intra-city DR

Service List

- Virtual Private Cloud (VPC)
- Elastic Cloud Server (ECS)
- RDS
- Data Replication Service (DRS)

Notes on Usage

- 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 MySQL**.

Prerequisites

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

2.4.2 Resource Planning

Table 2-13 Resource planning description

Category	Subcategor y	Planned Value	Remarks
RDS	RDS instance name	rds-mysql	Customize a name for easy identification.
	DB engine version	MySQL 5.7	-
	Instance type	Single	In this practice, select a single instance.
			To improve service reliability, selecting a primary/standby instance is recommended.
	Storage type	Cloud SSD	-
	AZ	AZ3	In this practice, select a single instance.
			To improve service reliability, create a primary/standby instance and then deploy them in two different AZs.
	Specification s	General-purpose 4 vCPUs 8 GB	-
DRS	Task name	DRS-mysql	Custom
migration task	Source DB engine	MySQL	In this practice, the source is a MySQL database built on an ECS.
	Destination DB engine	MySQL	In this practice, the destination is an RDS for MySQL instance.
	Network type	VPC	In this practice, select the VPC network.

2.4.3 Operation Process

The following figure shows the process of creating a MySQL database on an ECS, buying an RDS for MySQL instance, and migrating data from the MySQL database to the RDS instance.




Figure 2-27 Flowchart

2.4.4 Cloud Migration

2.4.4.1 Creating an RDS for MySQL Instance

Create an RDS for MySQL instance that is in the same VPC and security group as the self-managed MySQL database.

- Step 1 Go to the Buy DB Instance page.
- **Step 2** Configure basic information for the instance. Select **CN-Hong Kong** for **Region**.





Figure 2-28 Basic information

Step 3 Select an instance class and retain the default values for other parameters.

Figure 2-29 Instance class



Step 4 Click Next.

- **Step 5** Confirm the settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.

Step 6 Return to the instance list.

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

----End

2.4.4.2 Creating a Migration Task

This topic describes how to create a DRS migration task to migrate the **loadtest** database from the self-managed MySQL server to an RDS for MySQL instance.

Pre-migration Check

Before creating a migration task, check the migration environment to ensure smooth migration.

This example describes how to migrate data from a self-managed MySQL database to an RDS for MySQL instance. For more information, see **From MySQL to MySQL**.

Procedure

Migrate the **loadtest** database from a self-managed MySQL server to an RDS for MySQL instance.

- **Step 1** Go to the **Create Migration Task** page.
- **Step 2** Configure parameters as needed.
 - 1. Specify a migration task name. Select the region where the target instance is located, for example, **CN-Hong Kong**.

Figure 2-30 Migration task



2. Configure replication instance information.

Select the instance created in **Creating an RDS for MySQL Instance** as the destination instance.

Replication Instance Details ③ * Data Flow To the cloud Out of the cloud MySQL MongoDB MySQL schema and logic table Single-Node or Master/Standby Redis Redis cluster * Source DB Engine MySQL DDM TaurusDB * Destination DB Engine VPC

✓
⑤

✓
View DB Instance View Unselectable DB Instance + Destination DR Instance During the full synchronization of a DRS task, a lot or binlogs are generated. These binlogs may be temporarily stored locally, which may cause the storage space to be used up. You are advised to enable storage autoscaling for the RDS DB instance. During the DRS task, set an appropriate local referrition period for RDS binlogs. You can also clear binlogs exceeding the specified retention period with just a few close. * Replication Instance Subnet Full+Incremental Full Read-only Read/Write * Destination DB Instance Access During the mig be resumed. * Enable Binlog Cleanup **7** 0

Figure 2-31 Replication instance details

3. Select default for Enterprise Project.


Step 3 Click Create Now.

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

- **Step 4** Configure task information and click **Next**.
 - 1. Configure source database information.
 - 2. Click **Test Connection**.

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

Figure 2-32 Source database settings



- 3. Specify a username and password for the destination database.
- 4. Click **Test Connection**.

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




Figure 2-33 Destination database settings

Step 5 On the **Set Task** page, select the accounts and objects to be migrated, and click **Next**.

Select All for Migration Object.

For more information, see From MySQL to MySQL.

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

 If the check is complete and the check success rate is 100%, click **Next**.
- **Step 7** On the **Compare Parameters** page, click **Next** in the lower right corner to skip the comparison.
- Step 8 On the Confirm Task page, specify Start Time, Send Notifications, SMN Topic, Delay Threshold (s), and Stop Abnormal Tasks After, confirm that the configured information is correct, and click Submit to submit the task.
- **Step 9** After the task is submitted, view and manage it on the **Online Migration Management** page.

----End

2.4.4.3 Confirming Migration Results

You can check migration results with either of the following methods:

Automatic: **Viewing Migration Results on the DRS Console**. DRS automatically compares migration objects, users, and data of source and destination databases and provides migration results.


Manual: Viewing Migration Results on the RDS Console. You can log in to the destination instance to check whether the databases, tables, and data are migrated.

Viewing Migration Results on the DRS Console

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select **CN-Hong Kong**.

- **Step 3** Click the service list icon on the left and choose **Databases** > **Data Replication Service**.
- **Step 4** Click the name of the DRS migration task.
- **Step 5** In the navigation pane, choose **Migration Comparison**.
- **Step 6** By default, the **Object-Level Comparison** tab page is displayed. Click **Compare** and check the comparison results of the items such as databases, tables, and indexes between the source and destination databases.

Figure 2-34 Migration comparison



Step 7 Click the **Data-Level Comparison** tab, create a comparison task, and check the data comparison results between the source and destination databases.

If any check fails, rectify the fault by referring to Solutions to Failed Check Items.

----End

Viewing Migration Results on the RDS Console

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select **CN-Hong Kong**.
- Step 3 Click the service list icon on the left and choose Databases > Relational Database Service.
- **Step 4** Locate the required RDS instance and click **Log In** in the **Operation** column.
- **Step 5** In the displayed dialog box, enter the password and click **Test Connection**.
- **Step 6** After the connection test is successful, click **Log In**.
- **Step 7** Check and confirm the destination database name and table name. Check whether the data migration is complete.

----End

2.5 From ECS-hosted MySQL to TaurusDB

2.5.1 Overview

This practice describes how to install a MySQL database (community edition) on a Huawei Cloud ECS and create a TaurusDB instance, and use DRS to migrate data

from MySQL to TaurusDB. With DRS, you can perform real-time migration tasks with minimal downtime.

Scenarios

- With the rapid increase of enterprise workloads, traditional databases have poor scalability and require distributed reconstruction.
- Building traditional databases requires purchasing and installing servers, systems, databases, and other software. Its O&M is expensive and difficult.
- Traditional databases are poor in complex queries.
- It is hard for traditional databases to smoothly migrate data with no downtime.

Prerequisites


- You have created Huawei ID and completed real-name authentication.
- Your account balance is at least \$0 USD.

Solution Architecture

In this practice, the source database is a ECS-hosted MySQL instance and the destination database is a TaurusDB instance. **Figure 2-35** shows the deployment architecture when the ECS-hosted MySQL and TaurusDB instances are in the same VPC.

If the ECS-hosted MySQL and TaurusDB instances are not in the same VPC, you need to configure a **VPC peering connection** between the two VPCs. For details about the deployment architecture, see **Figure 2-36**.

Figure 2-35 Deployment architecture in the same VPC






Figure 2-36 Deployment architecture in different VPCs

TaurusDB Advantages

- Robust performance: GaussDB(for MySQL) decouples storage and compute and uses a "log as database" architecture and remote direct memory access (RDMA). It can deliver seven times the performance of open-source MySQL for certain service loads.
- Elastic scaling: In addition to a primary node, you can add up to 15 read replicas for a DB instance within minutes. You can also scale up or down CPU and memory specifications for a DB instance as needed.
- High reliability: DB instances can be deployed across AZs and there are three data copies under the shared distributed storage layer. A DB instance failover can be complete within seconds with a zero RPO.
- High security: With shared distributed storage, GaussDB(for MySQL) ensures zero data loss and service recovery within seconds. VPCs, security groups, SSL connections, and data encryption are used to strictly control secure access.
- High compatibility: GaussDB(for MySQL) is fully compatible with MySQL. You can easily migrate your MySQL databases to GaussDB(for MySQL) without reconstructing existing applications.
- Mass storage: Based on Huawei-developed data functions virtualization (DFV) distributed storage, GaussDB(for MySQL) supports up to 128 TB of storage.

Service List

- Virtual Private Cloud (VPC)
- Elastic Cloud Server (ECS)
- TaurusDB
- Data Replication Service (DRS)

Notes on Usage

The resources and test data in this practice are for demonstration only. Adjust them as needed.

For more information about TaurusDB data migration, see **From MySQL to TaurusDB**.

2.5.2 Prerequisites

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

2.5.3 Resource Planning

Table 2-14 Resource planning

Category	Subcategory	Planned Value	Remarks
VPC	VPC name	vpc-mysql	Customize a name for easy identification.
	Region	AP-Singapore	For low network latency and quick resource access, select the region nearest to you.
	AZ	AZ3	-
	Subnet	10.0.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-mysql	Customize a name for easy identification.
ECS (MySQL	ECS name	ecs-mysql	Customize a name for easy identification.
server)	Specifications	s6.xlarge.2 4 vCPUs 8 GiB	Select specification based on service requirements.
	OS	CentOS 7.6 64	-
	System disk	General purpose SSD 40 GiB	-
	Data disk	Ultra-high I/O, 100 GiB	-
	EIP	Auto assign	Buy an EIP because the public network is selected for the migration task.
ECS (MySQL client)	ECS name	ecs-client	Customize a name for easy identification.
	Specifications	s6.xlarge.2 4 vCPUs 8 GiB	Select specification based on service requirements.
	OS	CentOS 7.6 64	-

Category	Subcategory	Planned Value	Remarks
	System disk	General purpose SSD 40 GiB	-
	Data disk	Not required	1
	EIP	Auto assign	Buy an EIP as needed. If you do not need to access the client through a public network, you do not buy an EIP.
TaurusDB	Instance name	gauss-mysql	Customize a name for easy identification.
	DB engine	TaurusDB	•
	DB engine version	MySQL 8.0	-
	AZ type	Single-AZ	-
	AZ	AZ6	-
	Instance specifications	Dedicated Edition	-
	CPU architecture	x86 8 vCPUs 32 GB	-
DRS	Task name	DRS-TaurusDB	Customize a name.
migration task	Source DB engine	MySQL	In this example, take a MySQL instance (community edition) installed on an ECS as the source database.
	Destination DB engine	TaurusDB	In this example, take a TaurusDB instance as the destination database.
	Network type	Public	In this example, select the public network.

2.5.4 Operation Process

The process involves creating an ECS, installing a MySQL database on it, buying a TaurusDB instance, and migrating data from the database to the TaurusDB instance.

Figure 2-37 shows the main operation flowchart.




Figure 2-37 Flowchart

2.5.5 Procedure

2.5.5.1 ECS-hosted MySQL Server

2.5.5.1.1 Creating a VPC and Security Group

This section describes how to create a VPC and security group for your MySQL server and TaurusDB 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 AP-Singapore.

Step 3 Click in the upper left corner of the page and choose Networking > Virtual Private Cloud.

The VPC console is displayed.

- **Step 4** On the displayed page, click **Create VPC** in the upper right corner.
- **Step 5** Configure required parameters.



Step 6 Click Create Now.

Step 7 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 AP-Singapore.

Step 3 Click in the upper left corner of the page and choose Networking > Virtual Private Cloud.

The VPC console is displayed.

- **Step 4** In the navigation pane on the left, choose **Access Control** > **Security Groups**.
- **Step 5** Click **Create Security Group** in the upper right corner of the page.
- **Step 6** In the displayed dialog box, configure parameters as needed.
- Step 7 Click OK.
- **Step 8** Return to the security group list, locate the security group **sg-mysql**, and click its name.
- **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**.



Step 11 Perform Step 9 to Step 10 to allow access from database port 22.

After the rules were configured, the figure similar to the following is displayed.



----End

2.5.5.1.2 Creating an ECS (MySQL Server)


This section describes how to buy an ECS for installing a MySQL database (community edition).

- **Step 1** Log in to the management console.
- Step 2 Click in the upper left corner of the management console and select AP-Singapore.
- Step 3 Click in the upper left corner of the page and choose Compute > Elastic Cloud Server.


Step 4 Click Buy ECS.

Step 5 Configure ECS parameters.

 Set Specifications to General computing and select s6.xlarge.2 with 4 vCPUs and 8 GiB.




2. Select the image and disk specifications.




Step 6 Click **Next: Configure Network**.

 Select the VPC and security group created in Creating a VPC and Security Group.




2. Set **EIP** to **Auto assign**, **Billed by** to **Traffic**, and **Bandwidth Size** to **20**. The bandwidth size can be changed as required.



Step 7 Click **Next: Configure Advanced Settings**.

Specify ECS Name and Password.



- Step 8 Click Next: Confirm.
- **Step 9** Select an enterprise project, select the **Agreement** option, and click **Submit**.
- **Step 10** Return to the ECS list page and view the creation progress.

When the ECS status changes to **Running**, the ECS has been created.

----End

2.5.5.1.3 Installing a MySQL Database (Community Edition)

This section describes how to initialize disks and install a MySQL database (community edition).

Log In to the ECS

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select **AP-Singapore**.
- Step 3 Click in the upper left corner of the page and choose Compute > Elastic Cloud Server.
- **Step 4** Locate the ECS **ecs-mysql** and click **Remote Login** in the **Operation** column.
- **Step 5** Select **CloudShell-based Login**.
- **Step 6** Enter the password of user **root**.
 - **Ⅲ** NOTE

The password is the one you specified during the ECS creation.

----End

Initializing Disks

Step 1 Create the **mysql** folder.

mkdir /mysql

Step 2 View data disk information.

fdisk -l

The command output is as follows.

```
[root@ecs-mysql ~]# fdisk -l
Disk /dev/vda: 42.9 GB, 42949672960 bytes, 83886080 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk label type: dos
Disk identifier: 0x000e3a31
   Device Boot
                   Start
                                 End
                                          Blocks
                                                   Id System
/dev/vda1
                     2048
                            83886079
                                        41942016
                                                   83 Linux
Disk /dev/vdb: 107.4 GB, 107374182400 bytes, 209715200 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
```

Step 3 Initialize the data disk.

mkfs.ext4 /dev/vdb

Step 4 Attach the disk.

mount /dev/vdb /mysql

Step 5 Check whether the disk has been attached.

df -h

If the following output is returned, the disk has been attached.

```
[root@ecs-mysql ~]# df -h
               Size Used Avail Use% Mounted on
Filesystem
devtmpfs
               3.9G
                        0 3.9G
                                  0% /dev
tmpfs
               3.9G
                                 0% /dev/shm
                        0
                          3.9G
tmpfs
               3.9G
                     8.6M 3.9G
                                 1% /run
tmpfs
                                 0% /sys/fs/cgroup
               3.9G
                        0
                           3.9G
/dev/vda1
                40G
                     2.2G
                           36G
                                  6% /
tmpfs
               783M
                        0 783M
                                  0% /run/user/0
/dev/vdb
                99G
                      61M
                           94G
                                 1% /mysql
```

Step 6 Create a folder and switch to the **install** folder.

mkdir -p /mysql/install/data

mkdir -p /mysql/install/tmp

mkdir -p /mysql/install/file

mkdir -p /mysql/install/log

cd /mysql/install

- Step 7 Download and install the MySQL client.
- Step 8 Initialize the MySQL client.

/mysql/install/mysql-8.0.22/bin/mysqld --defaults-file= /etc/my.cnf -- initialize-insecure

Step 9 Start the MySQL client.

nohup /mysql/install/mysql-8.0.22/bin/mysqld --defaults-file= /etc/my.cnf &

Step 10 Connect to the MySQL client.

/mysql/install/mysql-8.0.22/bin/mysql

Step 11 Create user **root** and assign the required permissions to it.

grant all privileges on *.* to 'root'@'%' identified by 'xxx' with grant option;FLUSH PRIVILEGES;

----End

2.5.5.1.4 Creating an ECS and Installing a MySQL Client on It

Step 1 Create an ECS for installing a MySQL client by referring to **Creating an ECS** (MySQL Server).

■ NOTE

- This ECS must be in the same region, AZ, VPC, and security group as the ECS where the MySQL server is deployed.
- Data disks are not required.
- This ECS name is ecs-client.
- Other parameters are the same as those of the ECS where the MySQL server is deployed.
- **Step 2** Download and install the MySQL client. For details, see **How Can I Install the** MySQL Client?

----End

2.5.5.1.5 Testing the ECS-hosted MySQL Performance

- **Step 1** Log in to the management console.
- Step 2 Click in the upper left corner of the management console and select AP-Singapore.

- Step 3 Click in the upper left corner of the page and choose Compute > Elastic Cloud Server.
- **Step 4** In the ECS list, locate the ECS **ecs-client** and click **Remote Login** in the **Operation** column.
- **Step 5** On the displayed page, click **CloudShell**.
- **Step 6** Enter the password you specify during the ECS **ecs-client** creation.
- **Step 7** Download sysbench.

wget -c https://github.com/akopytov/sysbench/archive/1.0.18.zip

Step 8 Install sysbench.

unzip 1.0.18.zip

cd sysbench-1.0.18

./autogen.sh

./configure

make

make install

Step 9 Log in to the MySQL database and create test database **sbtest**.

mysql -u root -P 3306 -h < host> -p -e "create database sbtest"

Step 10 Import test background data to the **sbtest** database.

sysbench --db-driver=mysql --mysql-host=<host> --mysql-port=cort> --mysql-password> --mysql-db=sbtest -table_size=250000 --tables=25 --events=0 --time=600 oltp_read_write prepare

Step 11 Perform a pressure test.

sysbench --db-driver=mysql --mysql-host=<host> --mysql-port=<port> --mysql-user=<user> --mysql-password=<password> --mysql-db=sbtest -table_size=250000 --tables=25 --events=0 --time=600 --threads=<thread_num> --percentile=95 --report-interval=1 oltp_read_write run

NOTICE

Delete databases and tables to release storage space only after they are migrated using DRS.

sysbench --db-driver=mysql --mysql-host=<host> --mysql-port=<port> --mysql-password> --mysql-db=sbtest -- table size=250000 --tables=25 --events=0 --time=600 oltp read write cleanup

----End


2.5.5.2 Cloud Migration

This chapter describes how to create a TaurusDB instance, create a DRS migration task, and migrate data from the ECS-hosted MySQL server to the TaurusDB instance.


2.5.5.2.1 Creating a TaurusDB Instance

This section describes how to create a TaurusDB instance that is in the same VPC and security group as the ECS-hosted MySQL server.

- **Step 1** Log in to the management console.
- Step 2 Click in the upper left corner of the management console and select AP-Singapore.
- Step 3 Click in the upper left corner of the page and choose Databases > TaurusDB.
- Step 4 In the upper right corner, click Buy DB Instance.
- **Step 5** Configure the instance name and basic information.



Step 6 Configure instance specifications.



Step 7 Select the VPC and security group.


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



Step 8 Configure the instance password.



Step 9 Configure an enterprise project.



- Step 10 Click Next.
- **Step 11** After confirming the settings, click **Submit**.
- Step 12 Return to the instance list.

If the instance status becomes **Available**, the instance has been created.


----End

2.5.5.2.2 Creating a DRS Migration Task

This section describes how to create a DRS migration task to migrate the **sbtest** database from the ECS-hosted MySQL server to the TaurusDB instance.


- **Step 1** Log in to the management console.
- Step 2 Click in the upper left corner of the management console and select AP-Singapore.

- Step 3 Click in the upper left corner of the page and 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 task name.



2. Configure replication instance details as needed.

Set **Destination DB Instance** to the TaurusDB instance created in **Creating a TaurusDB Instance**.



3. Set Enterprise Project to default.



Step 6 Click Create Now.

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

- **Step 7** Configure source and destination database information.
 - 1. Configure source database information.
 - 2. Click Test Connection.

If a successful connection message is returned, you have logged in to the source database.

3. Configure the username and password for the destination database.

Click Test Connection.

If a successful connection message is returned, you have logged in to the destination database.

- Step 8 Click Next.
- **Step 9** Confirm the users, snapshots, and migration objects to be migrated.

Set Migrate Object to All.

- Step 10 Click Next.
- **Step 11** View pre-check results.
- **Step 12** If the check is complete and the check success rate is 100%, click **Next**.
- Step 13 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.

----End

2.5.5.2.3 Checking the Migration Results

You can check the migration results with either of the following methods:

Method 1: (Automatically) Check the migration results on the DRS console. DRS can compare migration objects, users, and data of source and destination databases and obtain the migration results.

Method 2: (Manually) Check the migration results on the TaurusDB console. Log in to the destination database to check whether the databases, tables, and data are migrated. Manually confirm the data migration status.

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 **AP-Singapore**.
- Step 3 Click in the upper left corner of the page and choose Databases > Data Replication Service.
- **Step 4** Click the DRS instance name.
- **Step 5** Click **Migration Comparison**.
- Step 6 Under the Compare Data Validate ALL Rows/Values and Compare Data Double Check During Cutover tabs, check whether the objects of the source database have been migrated to destination database.

----End

Checking the Migration Results on the TaurusDB Console

- **Step 1** Log in to the management console.
- Step 2 Click in the upper left corner of the management console and select AP-Singapore.
- Step 3 Click in the upper left corner of the page and choose Databases > TaurusDB.
- **Step 4** Locate the required TaurusDB instance and choose **More** > **Log In** in the **Operation** column.
- **Step 5** In the displayed dialog box, enter the password and click **Test Connection**.
- **Step 6** After the connection test is successful, click **Log In**.
- **Step 7** Check and confirm the destination database name and table name. Check whether the data migration is complete.

----End

Testing TaurusDB Performance

After the migration is complete, test TaurusDB performance by referring to **Performance White Paper**.

2.6 From ECS-hosted MongoDB 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:

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

Resource Planning

Table 2-15 Resource planning

Category	Subcategory	Planned Value	Description
VPC	VPC name	vpc-dds	Specify a name that is easy to identify.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	AZ	AZ1	-

Category	Subcategory	Planned Value	Description
	Subnet	10.0.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-default	Specify a name that is easy to identify.
ECS	ECS name	ecs-mongodb	Specify a name that is easy to identify.
	Specifications	s6.xlarge.2 4vCPUs 8GiB	The specifications in the left column are selected in this example.
			In your project, select specifications meeting your workload requirements. For details, see ECS Specifications .
	OS	CentOS 7.6 64	-
	System disk	General purpose SSD 40 GiB	-
	Data disk	Ultra-high I/O, 100 GiB	-
	EIP	Buy now	Buy an EIP because the public network is selected for the migration task.
DDS	DDS instance name	dds-test	Specify a name that is easy to identify.
	DB engine	DDS	-
	DB engine version	4.4	-
	AZ type	Single-AZ	-
	AZ	AZ1	-
	Instance specifications	Enhanced II	-
	CPU architecture	x86 8 vCPUs 32 GB	-
DRS migration task	Task name	DRS-dds	Specify a name that is easy to identify.
	Source DB engine	MongoDB	In this example, the source is a MongoDB database built on an ECS.
	Destination DB engine	DDS	In this example, a DDS instance serves as the destination database.
	Network type	Public network	Public network is used in this example.

Diagram

Figure 2-38 Source and destination databases in the same VPC






Figure 2-39 Source and destination databases in the same region and different VPCs



Migration Process

Figure 2-40 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*.

Preparations

1. Permissions:

Table 2-16 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 2-16 Required permissions

Database	Full Migration Permission	Full+Incremental Migration Permission
Source	 Replica set: The source database user must have the read permission for the 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 database sto be migrated and the config database. To migrate accounts and roles of 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 source database. To migrate accounts and roles of the source database user must have the read permission for the system.users and system.roles system tables of the admin database.

Database	Full Migration Permission	Full+Incremental Migration Permission
Destination	The destination database dbAdminAnyDatabase pedatabase and the readWidestination database. If the is a cluster instance, the have the read permission	ermission for the admin rite permission for the he destination database migration account must

Source database permissions:

The source MongoDB database user must have all the required permissions listed in **Table 2-16**. 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 Overview** in the *Virtual Private Cloud User Guide*.

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.

 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 2-17 Task settings

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.
Project	The project corresponds to the current region and can be changed.
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 2-18 Replication instance information

Parameter	Description
Data Flow	To the cloud
Source DB Engine	Select MongoDB.
Destination DB Engine	Select DDS .
Network Type	Select VPC .

Parameter	Description
Destination DB Instance	The DDS DB instance you purchased.
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.

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**.

DDS DB instance Source Database Type default_vpi C View VPC ? Subnet default_subnet-1-.123(IP Address or Domain Name ? Ensure that the entered addresses belong to the same DB instance. Authentication Database Database Username Database Password SSL Connection Test Connection **Destination Database** DB Instance Name dds-jiqun20-1 Database Username Database Password Test Connection

Figure 2-42 Source and destination database details Source Database

Table 2-19 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.

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 2-20 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 2-43 Migration object




Table 2-21 Migration object

Paramete r	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 <i>Data Replication Service User Guide</i> 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 $\mathbb C$ in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.
	- 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**.




Figure 2-44 Task Check

Ⅲ 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.

 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 2-45 Task startup settings




Table 2-22 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 Notification s	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 .
Synchroniza tion 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


- 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 2-46 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 2-47 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.
- 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()

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

2.7 From On-Premises MySQL to RDS for MySQL

DRS supports data migration from on-premises MySQL databases to RDS for MySQL instances. 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 for MySQL instance on the current cloud. The following network types are supported:

- Virtual Private Network (VPN)
- Public network

Diagram

Figure 2-48 VPN






Figure 2-49 Public network+SSL connection



Migration Process

Figure 2-50 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*.

Preparations

1. Permissions

Table 2-23 lists the permissions required for the source and destination databases when migrating data from on-premises MySQL databases to the RDS for MySQL instances on the current cloud.

Table 2-23 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, REFERENCES, and WITH GRANT OPTION.	
	If the destination databa 8.0.14 to 8.0.18, the SESS permission is required.	
	To migrate users, you mu INSERT, UPDATE, and DE MySQL database.	

Source database permissions:

The source database user must have all the required permissions listed in **Table 2-23**. 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 for 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 for 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 for MvSQL.

If the source database attempts to access the destination database through a public network, you do not need to configure the destination database.

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.

Figure 2-51 EIP of the replication instance



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.

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 for 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**.

Figure 2-52 Replication instance information



Table 2-24 Task settings

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.
Project	The project corresponds to the current region and can be changed.
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 2-25 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select MySQL.
Destination DB Engine	Select 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 RDS for MySQL instance you created.

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 remains 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**.




Figure 2-53 Source and destination database details

Table 2-26 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 2-27 Destination database settings

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

Parameter	Description
Database Username	The username for accessing the destination RDS for 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 2-54 Migration object

Table 2-28 Migration types and objects-public network

Paramete r	Description
Flow	You can choose whether to control the flow.
Control	 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 DATABAS E	During migration, executing DDL operations on the source database may affect the data migration performance to some extent. To reduce data migration risks, DRS allows you to filter out DDL operations.
	The database deletion operation can be filtered 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.

Paramete r	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. If you select Yes , you can select the accounts to be migrated as required.
	 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.

Figure 2-55 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 2-56 One-click modification





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.

Figure 2-57 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 *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 2-58 Task startup settings




Table 2-29 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.

Parameter	Description
Send Notification s	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 .
Synchroniza tion 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

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 2-59 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 2-60 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.

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. 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

2.8 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

Resource Planning

Table 2-30 Resource planning

Category	Subcategory	Planned Value	Description
VPC	VPC name	vpc-dds	Specify a name that is easy to identify.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	AZ	AZ1	-
	Subnet	10.0.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-default	Specify a name that is easy to identify.
On- premises MongoDB database	Database version	MongoDB 4.4	Specify a name that is easy to identify.
DDS	DDS instance name	dds-test	Specify a name that is easy to identify.
	DB engine	DDS	-
	DB engine version	4.4	-
	AZ type	Single-AZ	-
	AZ	AZ1	-
	Instance specifications	Enhanced II	-
	CPU architecture	x86 8 vCPUs 32 GB	-
DRS migration task	Task name	DRS-dds	Specify a name that is easy to identify.
	Source DB engine	MongoDB	In this example, the source is an on- premises MongoDB database.
	Destination DB engine	DDS	In this example, a DDS instance serves as the destination database.
	Network type	Public network	Public network is used in this example.

Diagram

Figure 2-61 VPN






Figure 2-62 Public network+SSL connection



Migration Process

Figure 2-63 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*.

Preparations

1. Permissions

Table 2-31 lists the permissions required for the source and destination databases when migrating data from on-premises MongoDB databases to DDS DB instances.

Table 2-31 Required permissions

Database	Full Migration Permission	Full+Incremental Migration Permission
Source	 Replica set: The source database user must have the read permission for the 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 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 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 source database. To migrate accounts and roles of the source database user must have the read permission for the system.users and system.roles system tables of the admin database.

Database	Full Migration Permission	Full+Incremental Migration Permission
Destination	The destination database user must have the dbAdminAnyDatabase permission for the adm database and the readWrite permission for the destination database. If the destination database is a cluster instance, the migration account muchave the read permission for the config database.	

- Source database permissions:

The source database user must have all the required permissions listed in **Table 2-31**. 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 2-64**.

Figure 2-64 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.

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.

- 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 2-65 Replication instance information



Table 2-32 Task settings

Parameter	Description		
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.		
Project	The project corresponds to the current region and can be changed.		
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 2-33 Replication instance settings

Parameter	Description		
Data Flow	Select To the cloud.		
Source DB Engine	Select MongoDB.		
Destination DB Engine	Select DDS .		
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 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. 		
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. 		

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 2-66 Source database information




Table 2-34 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 Huawei Cloud 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 2-67 Destination database information



Table 2-35 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 2-68 Migration object




Table 2-36 Migration object

Paramete r	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.	

Paramete r	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 $^{f C}$ 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.




Figure 2-69 Task Check

Ⅲ 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.

 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 2-70 Task startup settings




Table 2-37 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 Notification s	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.	
Synchroniza tion 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


- 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 2-71 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 2-72 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.
- 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()

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.

- 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

2.9 From RDS for MySQL to DDM

2.9.1 Overview

Description

This section describes:

- How to create an RDS for MySQL instance on Huawei Cloud.
- How to create a DDM instance on Huawei Cloud.
- How to migrate data from RDS for MySQL instances in different regions to DDM through a VPN.

Prerequisites

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


Service List

- Virtual Private Cloud (VPC)
- Virtual Private Network (VPN)
- Relational Database Service (RDS)
- Distributed Database Middleware (DDM)
- Data Replication Service (DRS)
- Data Admin Service (DAS)

Deployment Architecture

In this example, the source is a Huawei Cloud RDS for MySQL instance and the destination is a DDM instance in different regions of Huawei Cloud. Data is migrated from the source database to the destination database through a VPN. For details about the deployment architecture, see **Figure 2-73**.

Figure 2-73 VPN scenario



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.

2.9.2 Resource Planning

Table 2-38 Resource planning

Categor y	Subcategor y	Plan	Description
Source VPC	VPC name	vpc-DRSsrc	Specify a name that is easy to identify.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	AZ	AZ2	-
	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 databas	RDS instance name	rds-mysql-src	Specify a name that is easy to identify.
e)	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	DB engine version	MySQL 5.7	-
	Instance type	Single	A single instance is used in this example.
			To improve service reliability, select a primary/standby instance.
	Storage type	SSD	-
	AZ	AZ2	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.
	Specificatio ns	General-purpose 4 vCPUs 8 GB	-
Source VPN	Gateway name	vpngw-src	Specify a name that is easy to identify.

Categor y	Subcategor y	Plan	Description
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	Virtual Private Cloud (VPC)	vpc-DRSsrc	The value must be the same as the VPC of the source RDS DB instance.
	VPN connection name	vpn-src01	Specify a name that is easy to identify.
	Local Subnet	subnet-drs01	The value must be the same as the VPC subnet of the source RDS instance.
	Remote gateway	123.60.251.207	Specifies the peer VPN gateway. This parameter is set to the gateway address of the target VPN. After the target VPN is created, obtain the gateway information.
	Remote subnet	172.16.0.0/24	Specifies the peer VPN gateway. The value is the subnet of the destination VPN, which must be the same as the subnet of the VPC where the destination DDM instance resides.
Destinati on VPC	VPC name	vpc-DRStar	Specify a name that is easy to identify.
	Region	CN-Hong Kong	To achieve lower network latency, select the region nearest to you.
	AZ	AZ1	-
DDM (destinat ion databas e)	Subnet	172.16.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-drs02	Specify a name that is easy to identify.
	DDM instance name	ddm-drs-tar	Specify a name that is easy to identify.
	AZ	AZ1	You can select one or more AZs. You are advised to create the instance across different AZs to improve service reliability.

Categor y	Subcategor y	Plan	Description
	Node specification s	General-enhanced 8 vCPUs 16 GB	-
	Nodes	1	A single node has high availability risks. In practice, you are advised to create at least two nodes.
RDS instance associat	RDS instance name	rds-ddm01	Specify a name that is easy to identify.
ed with the DDM instance	Region	CN-Hong Kong	To achieve lower network latency, select the region nearest to you.
	DB engine version	MySQL 5.7	-
	Instance type	Single	A single instance is used in this example.
			To improve service reliability, select a primary/standby instance.
	Storage type	SSD	-
	AZ	AZ1	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.
	Specificatio ns	General-purpose 4 vCPUs 8 GB	-
Destinati on VPN	Gateway name	vpngw-tar	Specify a name that is easy to identify.
	Region	CN-Hong Kong	To achieve lower network latency, select the region nearest to you.
	Virtual Private Cloud (VPC)	vpc-DRStar	The value must be the same as the VPC of the destination DDM instance.

Categor y	Subcategor y	Plan	Description
	VPN connection name	vpn-tar01	Specify a name that is easy to identify.
	Local subnet	subnet-drs02	The value must be the same as the VPC subnet of the destination DDM instance.
	Remote gateway	123.60.236.84	Specifies the peer VPN gateway. In this example, this parameter is set to the gateway address of the source VPN. After the source VPN is created, the gateway address is obtained.
	Remote subnet	10.0.0.0/24	Specifies the peer VPN gateway. The value is the subnet of the source VPN, which must be the same as the subnet of the VPC where the source RDS for MySQL instance is located.
DRS migratio	Task name	DRS-MySQLToDDM	Specify a name that is easy to identify.
n task	Source DB engine	MySQL	In this example, the source is an RDS for MySQL instance on Huawei Cloud.
	Destination DB engine	DDM	In this example, the destination database is DDM.
	Network type	VPN	In this example, VPN is used.

2.9.3 Operation Process

Figure 2-74 shows the process of creating an RDS for MySQL instance and synchronizing data from an RDS for MySQL instance to DDM.




Figure 2-74 Flowchart

2.9.4 Preparing for the Source RDS for MySQL Instance

2.9.4.1 Creating a VPC and Security Group


Create a VPC and security group for an RDS 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 AP-Singapore.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

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 AP-Singapore.
- **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.



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




----End

2.9.4.2 Creating an RDS for MySQL Instance


Create an RDS for MySQL instance and construct test data.

Creating an Instance

- Step 1 Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select AP-Singapore.
- **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.



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** Confirm your settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Submit.
- **Step 11** Return to the instance list.


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

----End

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 AP-Singapore.
- **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 dialog box that is displayed, enter the username and password of the instance and click **Test Connection**.
- **Step 6** After the connection is successful, click **Log In** to log in to the RDS instance.
- **Step 7** Click **Create Database** to create the **db_test** database.



Step 8 Run the following statements in **db_test** to create table **table3**:



Step 9 Run the following statements in table **table3**_ to insert three lines of data:

```
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(1,'00:00:11','a');
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(2,'00:00:22','b');
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(5,'00:00:55','e');
```

```
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(1,'00:00:11','a')
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(2,'00:00:22','b')
INSERT INTO `db_test`.`table3_` (`Column1`,`Column2`,`Column3`) VALUES(5,'00:00:55','e')

Execute Back
```

----End


2.9.5 Preparing for the Destination DDM Instance

2.9.5.1 Creating a VPC and Security Group

Create a VPC and security group to prepare network resources and security groups for the destination DDM 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 CN-Hong Kong.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**. The VPC console is displayed.
- 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 CN-Hong Kong.
- 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.

Create Security Group

★ Name

sg-DR502

★ Template

General-purpose web server

The security group is for general-purpose web servers and includes default rules that allow all inbound ICMP traffic and inbound traffic on ports 22, 80, 443, and 3389. The security group is used for remote login, ping, and hosting a website on ECSs.

O/255

Show Default Rule


Cancel

Step 6 Specify the security group name and other information, and click **OK**.

----End

2.9.5.2 Creating a DDM Instance

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select CN-Hong Kong.
- Step 3 Click Service List on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the displayed page, in the upper right corner, click **Buy DDM Instance**.
- **Step 5** Specify the instance information and specifications as required.



Step 6 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 7** After the configuration is complete, click **Next** at the bottom of the page.
- **Step 8** Confirm your settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Submit.
- **Step 9** To view and manage the instance, go to the **Instances** page.

The default database port is **5066** and cab be changed after a DDM instance is created. If the status of the instance is **Running**, the instance has been created.

----End


2.9.5.3 Creating an RDS for MySQL Instance

Create an RDS for MySQL instance associated with the DDM instance.


Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select CN-Hong Kong.

- **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.



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**.



The RDS for MySQL instance must be in the same VPC and subnet as your DDM instance.



Step 8 Configure the instance password.



Step 9 Click Next.

- **Step 10** Confirm your settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.
- **Step 11** Return to the instance list.

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

----End


2.9.5.4 Creating a Schema and Associating It with the RDS 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 CN-Hong Kong.
- Step 3 Click Service List on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the **Instances** page, locate the required DDM instance and click **Create Schema** in the **Operation** column.
- **Step 5** On the displayed page, specify a sharding mode, enter a schema name, set the number of shards, select the required DDM accounts, and click **Next**.

In this example, the schema is unsharded, and the schema name is **db_test**.



Currently, only data can be migrated from the source RDS for MySQL database to the destination DDM instance. Source table structures and other objects cannot be migrated, so you need to create schemas in the destination database based on the table structure of the source database.




----End

2.9.5.5 Creating a DDM Account

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select CN-Hong Kong.
- **Step 3** Click **Service List** on the left and choose **Databases** > **Distributed Database Middleware**.
- **Step 4** On the **Instances** page, click the target instance name.
- **Step 5** In the navigation pane, choose **Accounts**.
- **Step 6** On the displayed page, click **Create Account** and configure the required parameters.

For details about the permissions required by the DDM account, see **Precautions**



Step 7 Click OK.

----End

2.9.5.6 Creating Table Structures in the Destination Database

Currently, only data can be migrated from the source RDS for MySQL database to the destination DDM instance. Source table structures and other objects cannot be migrated, so you need to create table structures and indexes in the destination database based on the table structures of the source schema. Objects that are not created in the destination database are not to be migrated. For more constraints, see **Before You Start**.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select CN-Hong Kong.
- Step 3 Click Service List on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the **Instances** page, locate the required instance and click **Log In** in the **Operation** column.
- **Step 5** On the displayed page, enter the username and password created in **Creating a DDM Account**, and then click **Test Connection**.
- **Step 6** After the connection is successful, click **Log In** to log in to the DDM instance.

- Step 7 Click the db_test schema created in Creating a Schema and Associating It with the RDS for MySQL Instance.
- **Step 8** Run the following statements in database **db_test** to create table **table3**_ with the same structure as the source table:

```
CREATE TABLE 'db_test'.'table3_' (

'Column1' INT(11) UNSIGNED NOT NULL,

'Column2' TIME NULL,

PRIMARY KEY ('Column1')

ENGINE = InnoDB

DEFAULT CHARACTER SET = utf8mb4

COLLATE = utf8mb4_general_ci;
```

----End

2.9.6 Creating and Configuring a VPN

2.9.6.1 Creating a VPN for the Source Database


- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner of the management console and select AP-Singapore.
- Step 3 Click Service List in the upper left corner. Under Networking, select Virtual Private Network.
- **Step 4** In the navigation pane on the left, choose **Virtual Private Network** > **VPN Gateway**.
- **Step 5** On the **VPN Gateways** page, click **Buy VPN Gateway**.
- **Step 6** Configure the required parameters.



A CAUTION


The VPC connected to the VPN gateway must be the same as the VPC of the source RDS for MySQL instance, that is, the VPC created in **Creating a VPC and Security Group**.

Step 7 Specify the VPN connection information as prompted and click Buy Now.



CAUTION


- The local subnet must be the same as the subnet of the VPC where the source RDS for MySQL DB instance is located, that is, the subnet created in Creating a VPC and Security Group.
- The remote gateway and remote subnet are the gateway and subnet of the target VPN. The target VPN has not been created. You can enter any value and change it after the target VPN is created.
- **Step 8** After the VPN gateway is created, view its information in the VPN gateway list. The status of the VPN gateway is **Not connected**. When a VPN connection uses this VPN gateway, the VPN gateway enters the **Normal** state.



----End

2.9.6.2 Creating a VPN for the Destination Database


- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select CN-Hong Kong.
- Step 3 Click Service List in the upper left corner. Under Networking, select Virtual Private Network.
- **Step 4** In the navigation pane on the left, choose **Virtual Private Network > VPN Gateway**.
- **Step 5** On the **VPN Gateways** page, click **Buy VPN Gateway**.
- **Step 6** Configure the required parameters.



CAUTION

The VPC connected to the VPN gateway must be the same as the VPC of the destination DDM instance, that is, the VPC created in **Creating a VPC and Security Group**.

Step 7 Specify the VPN connection information as prompted.



A CAUTION

- The local subnet must be the same as the subnet of the VPC where the
 destination DDM instance is located, that is, the subnet created in Creating a
 VPC and Security Group.
- The remote gateway and remote subnet are the gateway and subnet of the source VPN. Configure the parameters based on the information about the VPN created in Creating a VPN for the Source Database.

----End

2.9.6.3 Modifying the VPN Configuration for the Source Database

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select AP-Singapore.
- Step 3 Click Service List in the upper left corner. Under Networking, select Virtual Private Network.
- **Step 4** In the navigation pane on the left, choose **Virtual Private Network > VPN Connections**.
- **Step 5** On the **VPN Connections** page, locate the row that contains the VPN connection created in **Creating a VPN for the Source Database** and click **Modify** in the **Operation** column.

Step 6 On the **Modify VPN Connection** page, change the values of **Remote Gateway** and **Remote Subnet**.





The remote gateway and remote subnet are the gateway and subnet of the destination VPN. Configure the parameters based on the information about the VPN created in Creating a VPN for the Destination Database.

Step 7 After the configuration is complete, view the VPN gateway information in the list. The VPN gateway status is **Normal**.



----End

2.9.7 Creating a DRS Migration Task

Create a DRS migration task to migrate data from RDS for MySQL databases in different regions to DDM.

Pre-Migration Check


Before creating a migration task, check the migration environment.

This section describes how to migrate data from a MySQL database to DDM. 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 CN-Hong Kong.
- **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 the required parameters.
 - 1. Specify a migration task name.



2. Configure replication instance details as needed.

Select the DDM instance created in **Creating a DDM Instance** as the destination database.

Figure 2-75 Replication instance information



Step 6 Click Create Now.

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

- **Step 7** Configure source and destination database information.
 - 1. Configure the source database information and click **Test Connection**. If a successful test message is returned, the database is connected.



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

Destination Database



Step 8 Click Next.

- **Step 9** On the **Set Task** page, configure migration objects.
 - Migrate Object: Select Tables.

Step 10 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 11 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 **Full migration**, the migration task has been started.

- For migration from MySQL to DDM, full migration and full+incremental migration modes are supported.
- 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

2.9.8 Confirming Migration Results

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

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


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

Viewing 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 CN-Hong Kong.
- **Step 3** Under the service list, choose **Databases** > **Data Replication Service**.
- **Step 4** Click the DRS instance name.
- **Step 5** In the navigation pane on the left, choose **Migration Comparison**.
- **Step 6** Click the **Object-Level Comparison** tab and check whether some objects are missing.
 - Click **Compare**. After the comparison is complete, view the comparison results.



- **Step 7** Choose **Data-Level Comparison** and check whether the number of rows of migrated objects is consistent.
 - 1. Click Create Comparison Task.
 - 2. In the displayed dialog box, select the comparison type, time, and object.



3. After the comparison task is complete, view the data comparison results.



To view the comparison details, click View Results next to the comparison task.



----End

Viewing Migration Results on the DDM Console

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select CN-Hong Kong.
- Step 3 Click Service List on the left and choose Databases > Distributed Database Middleware.
- **Step 4** Locate the DDM 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 databases and tables of the source instance have been migrated.

----End

2.10 From MySQL Schema and Logic Table to DDM

2.10.1 Overview

Scenarios


Traditional databases are expensive and difficult to maintain. These deficiencies are some of the driving forces behind cloud migration. This section describes how

to use DRS to migrate MySQL shards and tables to a DDM instance on Huawei Cloud.

Solution Architecture

In this example, DDS shards and tables are associated with two MySQL instances at the source side, and a DDM instance is associated with two RDS MySQL instances at the destination side. Source data is migrated to the destination through a DRS instance, and source data is continuously synchronized before the incremental synchronization task is manually stopped. Figure 2-76 shows the overall deployment architecture.

Figure 2-76 Solution architecture



Advantages

- The full+incremental migration mode is used to ensure that services are not interrupted and data is not lost during data migration.
- DRS saves traditional database administrator (DBA) costs and hardware costs, and supports on-demand pricing.
- A migration task can be set up within minutes.

2.10.2 Resource Planning

Table 2-39 Resource planning

Resource	Name	Description	Quant ity
Virtual Private Cloud (VPC)	vpc-src-172	CIDR: 172.16.0.0/16	1
Source DDM instance	Autotest- DDM- SRC-001	 Specifications: 4 vCPUs 8 GB Nodes: 2 Schema: db_test_info Database user: test_info 	1

Resource	Name	Description	Quant ity
RDS DB instance associated with the source DDM instance	Auto-ddm- Single- SRC-001	 DB instance type: single Database version: MySQL 5.7 Specifications: General-enhanced 2 vCPUs 4 GB Storage type: ultra-high I/O 	1
	Auto-ddm- Single- SRC-002	 DB instance type: single Database version: MySQL 5.7 Specifications: General-enhanced 1 vCPU 2 GB Storage type: ultra-high I/O 	1
Destination DDM instance	Autotest- DDM- SRC-001	 Specifications: 2 vCPUs 8 GB Nodes: 2 Schema: db_test_info Database user: test_info_tar 	1
RDS DB instance associated with the destination DDM	Auto-tar- ddm-0001	 DB instance type: single Database version: MySQL 5.7 Specifications: General-enhanced 1 vCPU 2 GB Storage type: ultra-high I/O 	1
instance	Auto-tar- ddm-0002	 DB instance type: single Database version: MySQL 5.7 Specifications: General-enhanced 1 vCPU 2 GB Storage type: ultra-high I/O 	1
DRS replication instance	DRS-test- info	 Source DB engine: MySQL shards and tables Destination DB engine: DDM Network type: public network 	1

2.10.3 Process

Figure 2-77 shows the main operation flowchart.

Figure 2-77 Flowchart




Table 2-40 Migration process

No.	Procedur e	Description
1	Construct ing Data Before the Migratio n	Before the migration, construct data in the source database so that you can check whether the migration is successful after the migration.
2	Preparing for the Destinati on DDM Instance	 Creating a VPC and Security Group: Prepare network resources for creating an instance. Creating a DDM Instance: Create a destination DDM
		instance.
		3. Creating an RDS for MySQL Instance: Create an RDS for MySQL instance associated with the DDM instance.
		4. Creating a DDM Account: Create an account for connecting to the DDM schema.
		 Creating a Schema and Associating It with the RDS for MySQL Instance: Associate DDM with an RDS for MySQL instance.
		6. Creating Table Structures in the Destination Database: Create a table structure in the destination database for data migration.

No.	Procedur e	Description
3	Migratin g Database s	Describe how to create a DRS migration task and the overall migration process.
4	Verifying Data After Migratio n	Verify the data migration results by comparing data.

2.10.4 Procedure

2.10.4.1 Constructing Data Before the Migration

Before the migration, you need to construct some data types in the source database for data verification after the migration.

Prerequisites

You have obtained the IP address, port, username, and password of the source instance.

Procedure

- **Step 1** Use the database connection tool to connect to the IP address of the source database.
- **Step 2** Construct data in the source database based on the supported data types.
 - 1. Log in to the source database.
 - 2. Open a schema to be migrated.
 - The schema used in this practice is **db_test_info**.
 - 3. Clear the previous table information in the **db_test_info** schema to ensure that data can be created successfully.
 - drop table if exists db_test_info_001;
 - **db_test_info_001** is the name of the table created in this practice. Replace it as required.
 - 4. Create a data table.

```
create table db_test_info_001(
column19 TIMESTAMP,
column18 VARCHAR(256),
column11 DATETIME,
column10 ENUM('a','b','c'),
```

```
column13 FLOAT(2,1),
 column12 YEAR,
 column15 VARCHAR(64) not null,
 column14 DOUBLE(2,1),
 column17 INT,
 column16 LONGTEXT,
 column20 TIME,
 column21 BIT(8),
 column22 LONGBLOB,
 column23 MEDIUMINT,
 column24 VARCHAR(20),
 column9 DECIMAL(2,1),
 column8 VARCHAR(1024),
 column5 SMALLINT,
 column4 DATE,
 column7 VARCHAR(32),
 column6 SET('hehe','xixi','haha'),
 column1 char(1),
 column0 MEDIUMINT,
 column3 BIGINT,
 column2 TINYINT,
 primary key(column19,column11,column17)
 );
Insert records into the table.
 insert into
 db test info 001(column19,column11,column10,column1,column24,colu
 mn15,column23,column0,column8) values('2019-07-24
 14:08:58',now(),'A',(RAND()*10000000),
 (RAND()*10000000),'db_test_info_001',(RAND()*10000000),
 (RAND()*1000000),'a');
 insert into
 db test info 001(column19,column11,column17,column10,column1,colu
 mn24,column15,column23,column0,column8) values(now(),now(),
 (RAND()*10000000),'A',(RAND()*10000000),
 (RAND()*10000000),'db_test_info_001',(RAND()*10000000),
 (RAND()*1000000),'a');
 insert into
 db test info 001(column19,column11,column17,column10,column1,colu
 mn24,column15,column23,column0,column8) values(now(),now(),
 (RAND()*10000000),'A',(RAND()*10000000),
 (RAND()*10000000),'db_test_info_001',(RAND()*10000000),
 (RAND()*1000000),'a');
 insert into
 db test info 001(column19,column11,column17,column10,column1,colu
 mn24,column15,column23,column0,column8) values(now(),now(),
```

```
(RAND()*10000000),'A',(RAND()*10000000),
(RAND()*10000000),'db_test_info_001',(RAND()*10000000),
(RAND()*1000000),'a');
```

6. Query the results.

select * FROM db test info 001;

----End

2.10.4.2 Preparing for the Destination DDM Instance

2.10.4.2.1 Creating a VPC and Security Group

This section describes how to create a VPC and security group for the instance you will create.

Prerequisites


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


Creating a VPC

- **Step 1** Log in to the **Huawei Cloud console**.
- **Step 2** Click on the upper left corner and select a region and a project.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

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 **Huawei Cloud console**.
- **Step 2** Click on the upper left corner and select a region and a project.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**. The VPC console is displayed.
- **Step 4** 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 security group name **sg-01**.
- Step 9 Click the Inbound Rules tab, and then click Add Rule.



Step 10 Configure an inbound rule and add the IP address of the source database.




----End

2.10.4.2.2 Creating a DDM Instance

This section describes how to create a DDM instance as the destination database for the migration task.

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** Click on the upper left corner and select a region and a project.
- Step 3 Click the service list icon on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the displayed page, in the upper right corner, click **Buy DDM Instance**.
- **Step 5** Specify a node class and other information.



Step 6 Select a VPC and security group for the DB instance and configure the database port.

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

- **Step 7** After the configuration is complete, click **Next** at the bottom of the page.
- **Step 8** Confirm your settings.
 - To modify your settings, click Previous.
 - If you do not need to modify your settings, click **Submit**.
- **Step 9** Go to the **Instances** page to view and manage the instance.

The default database port is **5066** and cab be changed after the instance is created. If the status of the instance is **Running**, the instance is created.


----End

2.10.4.2.3 Creating an RDS for MySQL Instance


Create an RDS for MySQL instance and associate it with the DDM instance.

Procedure

- **Step 1** Click on the upper left corner and select a region.
- **Step 2** Under the service list, choose **Databases** > **Relational Database Service**.
- Step 3 Click Buy DB Instance.
- **Step 4** Configure the instance name and basic information.



Step 5 Configure instance specifications.



Step 6 Select a VPC and security group for the DB instance and configure the database port.

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



The RDS for MySQL instance must be in the same VPC and subnet as the DDM instance.

Step 7 Configure the instance password.



Step 8 Click Next.

- **Step 9** Confirm your settings.
 - To modify your settings, click Previous.
 - If you do not need to modify your settings, click Submit.
- **Step 10** Return to the instance list.

If the instance status becomes **Available**, the instance is created.


Create another RDS DB instance based on the destination RDS DB instance information in **Resource Planning**.

----End

2.10.4.2.4 Creating a DDM Account

- **Step 1** Click on the upper left corner and select a region.
- **Step 2** Click **Service List** on the left and choose **Databases** > **Distributed Database Middleware**.
- **Step 3** On the **Instances** page, locate the required DDM instance and click its name.
- **Step 4** In the navigation pane, choose **Accounts**.
- **Step 5** On the displayed page, click **Create Account** and configure parameters as needed.

For details about the permissions required by the DDM account, see **Precautions**. In this practice, select all permissions.



Step 6 Click OK.

----End


2.10.4.2.5 Creating a Schema and Associating It with the RDS for MySQL Instance

- **Step 1** Click in the upper left corner and select a region.
- Step 2 Click Service List on the left and choose Databases > Distributed Database Middleware.
- **Step 3** On the **Instances** page, locate the required DDM instance and click **Create Schema** in the **Operation** column.
- **Step 4** On the displayed page, specify a sharding method, enter a schema name, set the number of shards, select the required DDM accounts, and click **Next**.

In this example, the schema is unsharded, and the schema name is **db_test_info**.



Currently, only data can be migrated from the RDS for MySQL to DDM. To migrate table structures and other objects, you need to create schemas in the destination DDM instance based on table structures of the source RDS for MySQL instance.



----End

2.10.4.2.6 Creating Table Structures in the Destination Database

Currently, DRS can only migrate data from the source instance to the destination DDM instance. It cannot migrate table structures or other objects in the source database. To migrate table structures and other objects, you need to create table structures and indexes in the destination database based on table structures of the source schema. Any source objects that have no corresponding objects created in the destination cannot be migrated. For more constraints, see **Before You Start**.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click oin the upper left corner and select a region.
- **Step 3** Click **Service List** on the left and choose **Databases** > **Distributed Database Middleware**.
- **Step 4** On the **Instances** page, locate the required instance and click **Log In** in the **Operation** column.
- **Step 5** On the displayed page, enter the username and password created in **Creating a DDM Account**, and click **Test Connection**.

- **Step 6** After the connection is successful, click **Log In** to log in to the DDM instance.
- Step 7 Click the db_test_info schema created in Creating a Schema and Associating It with the RDS for MySQL Instance.
- **Step 8** Run the following SQL statements in database **db_test_info** to create table **db_test_info_001** with the same structure as the source:

```
create table db_test_info_001(
column19 TIMESTAMP,
column18 VARCHAR(256),
column11 DATETIME,
column10 ENUM('a','b','c'),
column13 FLOAT(2,1),
column12 YEAR,
column15 VARCHAR(64) not null,
column14 DOUBLE(2,1),
column17 INT,
column16 LONGTEXT,
column20 TIME,
column21 BIT(8),
column22 LONGBLOB,
column23 MEDIUMINT,
column24 VARCHAR(20),
column9 DECIMAL(2,1),
column8 VARCHAR(1024),
column5 SMALLINT,
column4 DATE,
column7 VARCHAR(32),
column6 SET('hehe','xixi','haha'),
column1 char(1),
column0 MEDIUMINT,
column3 BIGINT,
column2 TINYINT,
primary key(column19,column11,column17)
);
----End
```

2.10.4.3 Migrating Databases

Create a DRS instance and migrate table data from the source RDS for MySQL database to the **db_test_info** schema of the DDM instance.

Pre-migration Check


Before creating a migration task, you need to manually check the migration conditions for smooth migration.

Before the migration, see **Before You Start**.


Creating a Migration Task

- **Step 1** Log in to the **Huawei Cloud console**.
- **Step 2** Click in the upper left corner and select a region.

 Select the region where the destination instance is deployed.
- **Step 3** Click the service list icon on the left and choose **Databases** > **Data Replication Service**.
- **Step 4** In the navigation pane on the left, choose **Online Migration Management** and click **Create Migration Task**.
- **Step 5** Configure the replication instance information.
 - 1. Select a region and project, and enter a task name.




2. Configure the migration task type and select the destination instance and subnet.




- Click Create Now.
- **Step 6** Configure the source and destination instance information.
 - 1. Enter the IP address, port number, username, and password of the source database.

Click Test Connection.




Enter the username and password of the destination database. Click **Test Connection**.

Destination Database




3. Click **Next**. In the displayed box, read the message carefully and click **Agree**.




Step 7 Configure the synchronization task.

Select the databases and tables of the source database to be migrated. In this practice, select the **db_test_info_001** table in database **db_test_info**.




Step 8 Perform a pre-check.

1. After all settings are complete, perform a pre-check to ensure that the migration is successful.




2. If any check fails, review the cause and rectify the fault. Then, click **Check Again**.



If all check items pass the pre-check, click Next.

Step 9 Confirm the task.

1. Check whether all configured information is correct.



- 2. Specify **Start Time** and select the check box before the agreement.
- 3. Click Submit.

×

Notice




When DRS migration tasks are in progress, certain operations may cause issues. To ensure migration success, we strongly recommend that you read the migration precautions carefully before starting migration tasks and follow the instructions to ensure migration stability.



Any task that is active will be billed, even if its status becomes abnormal. If a task is no longer needed, stop the task to avoid unnecessary fees.


If the task status is abnormal for more than 14 days, the task automatically stops. Pay attention to the alarms you received and handle the task in time to resume the download and avoid task retry failure.



Start at a specified time

Step 10 After the task is submitted, view and manage it.

After the task is created, return to the task list to view the status of the created task.



----End


2.10.4.4 Verifying Data After Migration

When the task status changes to **Incremental migration**, the full migration is complete. Then, you can log in to the DDM instance to view the migration results.


Step 1 Wait until the synchronization task status becomes **Incremental**.



- **Step 2** Click the task name to go to the **Basic Information** page.
- **Step 3** Verify data consistency.
 - Choose Migration Comparison > Object-Level Comparison to view the database and table migration results.



2. Choose **Migration Comparison** > **Data-Level Comparison**, click **Create Comparison Task**, and view the migration results of the rows in the table.




Step 4 Stop the migration task.

After data is completely migrated to the destination database, stop the migration task.

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



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



Step 5 After the migration is complete, test the performance.

For details about how to test the DDM performance, see **DDM Performance White Paper**.

----End

3 Backup Migration

3.1 Migrating Microsoft SQL Server Backup Data to RDS for SQL Server

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

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

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

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 for 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		

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.




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.





Figure 3-6 Configuring compression parameter

Step 3 Export the full backup file.

- 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**.

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**.




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**.

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**.




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.

■ 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, download and install the OBS Browser+ client. For details, see Object
 Storage Service Client Guide.

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. 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. For details, see **Logging 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



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**.

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**.

Decination AGD SB Instance Name

Education File Format

Education File Format

Education File Format

Education File Interveneral
Full Interveneral
Location PreSection Address Involved Fremain available, but incremental backups can no longer be restored.

Cherwrite Data

No
The destination databases involved remain available, but incremental backups can no longer be restored.

Cherwrite Data
No
The destination databases involved remain available, but incremental backups can no longer be restored.

Cherwrite Data
No
The destination of databases involved remain available, but incremental backups can no longer be restored.

Perform Pre-venification
No
Restore Database

Ad
Custom
①

So

So

Solved Database Name

O

Solved Database Name

Solve

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.
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.

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**.

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

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

Migrate the full backup. Select No for Last Backup File Step 1 full bak Full backup Step 2 restore the backup file : Step N Select No for Last Backup File N transaction log backup Migrate the last transaction log Select Yes for Last Backup File.

Figure 3-15 Full+incremental migration

Migration Process

Check the local database Export the full backup of the Export the (incremental) log backup file. Upload the backup file to the specified OBS bucket. Log in to the DRS console and create a backup migration task. Continue to perform incremental migration. Perform the migration

Figure 3-16 Flowchart

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 on 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 .

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.
Backup Type	Select Full. Full: indicates full backup files.

Description
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.
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.
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.
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

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**.

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**.




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**.

Database Information

* Destination RDS DB Instance Name

* Backup Type

Full Incremental
Incremental: indicates transaction log backup files.

* Last Backup File ②

Yes

No

If you select No, the destination databases involved in this migration will be unavailable and incremental backups can still be restored.

* Perform Pre-verification

Yes

No

* Restore Database

All Custom

①

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 vestered.
	• 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.

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.

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**.

Backup File Information Microsoft SQL Server * Database Type * Backup File Source OBS Bucket RDS full backup ? drs ▼ C ③ * Bucket Name

Figure 3-21 Incremental backup

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.

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**.

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:

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




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.

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**.

3.1.7 Manual Configuration

Scenarios

After data is migrated from the local host or VMs to the RDS for 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 for 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

password as those of your local Login account on the RDS for 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

IF (SUSER_ID('sa') IS NULL) EBGIN CREATE LOGIN [sa] WITH PASSWORD = 0x0100396F2EFAD6A3D84E2AEB941E8ED32E5189A48E757...

IF (SUSER_ID('rdsuser2') IS NULL) EBGIN CREATE LOGIN [rdsuser2] WITH PASSWORD = 0x01008E86CBC25FC67008D4EE75AD660D1...

IF (SUSER_ID('csidbo') IS NULL) EBGIN CREATE LOGIN [csidbo] WITH PASSWORD = 0x0100A508789C15CE6880648162A5EDF4F4D2E...

IF (SUSER_ID('rdsuser3') IS NULL) EBGIN CREATE LOGIN [rdsuser3] WITH PASSWORD = 0x010073DA9A7986677E8AF7077EF67...

IF (SUSER_ID('rdsuser3') IS NULL) EBGIN CREATE LOGIN [rdsuser3] WITH PASSWORD = 0x01009448FEDDE28ED5E529384028CA0...

6 IF (SUSER_ID('rest2') IS NULL) EBGIN CREATE LOGIN [Test2] WITH PASSWORD = 0x0100130953CEEAEC997D0886BAF65F848BCAA44...

7 IF (SUSER_ID('rest3') IS NULL) EBGIN CREATE LOGIN [Test3] WITH PASSWORD = 0x01000EE98873948E02599EDCD9538426637281B7...

8 IF (SUSER_ID('Test4') IS NULL) EBGIN CREATE LOGIN [Test4] WITH PASSWORD = 0x010006E98873948E02599EDCD9538426637281B7...

9 IF (SUSER_ID('Test5') IS NULL) EBGIN CREATE LOGIN [Test5] WITH PASSWORD = 0x0100056EBF845DF098D2DF9395AF7B7018A20735...
```

- **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 for SQL Server DB instance to ensure permission consistency.

```
declare @DBName 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
```

```
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.




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.

□ 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.

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'

■ NOTE

The owner account of the new job is very important. On the RDS for 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 for 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]
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb1', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb1.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb2', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb2.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb3', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb3.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb4', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb4.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb5', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb5.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb6', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb6.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
ALTER DATABASE [tempdb] ADD FILE ( NAME = N'tempdb7', FILENAME =
N'D:\RDSDBDATA\Temp\tempdb7.ndf', SIZE = 65536KB, FILEGROWTH = 65536KB)
GO
```




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 for 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 for 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 for SQL Server instance to the same.

| DifferentialBarel M NUMERIC CS. 01
| Distance Enquelos
| Distance Enquelos
| Distance Enquelos
| Distance Enquelos
| DifferentialBarel M NUMERIC CS. 01
| Distance Enquelos
| Di

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



4 Real-Time Synchronization

4.1 From Other Cloud PostgreSQL to RDS for 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 PostgreSQL databases






Figure 4-2 PostgreSQL databases on other cloud servers

Synchronization Process

Start Create a Synchronization Task. Check the Synchronization Progress. Compare Synchronization Items. Synchronize services. End

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.

• 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	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, and the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords)	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords), the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections NOTE • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: 1. Add host replication connections, perform the following steps: 1. Add host replication connections, perform the following steps: 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Synchronization Type	Full	Full+Incremental
Destination	Database-level: The Orequired.	CREATEDB permission is
	Table-level:	
	 To synchronize data permission is require 	abases, the CREATEDB red.
	and CREATE permis	hema, the CONNECT ssions for the database chema are required.
	contains the schem	on for the database that a, and the USAGE and s for the schema that
	 Synchronization user: permission is required 	
	privilege cannot be mo object permissions of	permissions: The default odified. Otherwise, the the destination database with those of the source
		ers, text search parsers, and destination database version QL 11.11 or later, and the must be user root or a

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.

- 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.

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, as shown in **Figure 4-4**.

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.
 - Log in to the management console and choose Databases > Data Replication Service to go to the DRS console.
 - 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 the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.
Project	The project corresponds to the current region and can be changed.
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 RDS for PostgreSQL instance you created.	
Synchronization Mode	 Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. 	
	 Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. 	
	– Incremental	
	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 remains 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.	
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 for PostgreSQL instance you have selected during the synchronization instance creation is displayed by default and cannot be changed.

Parameter	Description
Database Username	The username for accessing the destination RDS for PostgreSQL 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.

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.
Synchronizati on 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 objects are selected before synchronization or select database-level synchronization.


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.

Figure 4-8 Task Check



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

Figure 4-9 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 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 ECS-hosted PostgreSQL to RDS for 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 4-10 Source and destination databases in the same VPC






Figure 4-11 Source and destination databases in the same region and different VPCs



Synchronization Process

Figure 4-12 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*.

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	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, and the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords)	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords), the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections NOTE • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: 1. Add host replication connections, perform the following steps: 1. Add host replication connections, perform the following steps: 2. Run select pg_reload_conf(); in the source database. 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Synchronization Type	Full	Full+Incremental
Destination	 Database-level: The CREATEDB permission is required. Table-level: 	
		abases, the CREATEDB red.
	and CREATE permis	hema, the CONNECT ssions for the database chema are required.
	contains the schem	on for the database that a, and the USAGE and s for the schema that
	 Synchronization user: permission is required. 	
	privilege cannot be mo object permissions of	permissions: The default odified. Otherwise, the the destination database with those of the source
	NOTE To synchronize event trigge text search templates, the comust be RDS for PostgreSQ destination database user member of user root.	destination database version L 11.11 or later, and the

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.

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 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

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.

Syndromization instance Ceptals So
The Manager and a series of a read of the reg bit in entires

- Code from

- Notice for the Share of the reg bit in entires

- Code from

- Notice for the Share of the reg bit in entires

- Share of Expert

- Notice for the Share of Expert

-

Figure 4-13 Synchronization instance details

Table 4-8 Task settings

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.
Project	The project corresponds to the current region and can be changed.
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: =<>&'\"

Parameter Description Data Flow To the cloud Source DB Select **PostgreSQL**. Engine **Destination DB** Select **PostgreSQL**. Engine 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 The PostgreSQL instance you purchased. Instance Synchronization Full+Incremental Type 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. All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. Incremental 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.

Table 4-9 Synchronization instance settings

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**.

□ 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 4-14 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.
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

Source Database

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

Source Database Type

Self-butt on ECS

ROS DN Instance

DI Instance Name

Ind
User DB Instance View Usrelectable DB Instance

Database Dsername

Database Password

......

Figure 4-15 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 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 4-16 Destination database settings

Test Connection



Table 4-12 Destination database settings

Parameter	Description
DB Instance Name	The RDS for 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.

Textification for the control of the

Figure 4-17 Synchronization mode

Table 4-13 Synchronization object

Parameter	Parameter Description	
raiailletei	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.	

Parameter	Description
Synchronizati on 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 objects are selected before synchronization or select database-level synchronization.
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.




Figure 4-18 Task Check

□ NOTE

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

7. On the Confirm Task page, specify Start Time and click Next.

- 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

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

Figure 4-19 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

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 for PostgreSQL


DRS supports data synchronization from on-premises MySQL databases to RDS for PostgreSQL instances. 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 an RDS for PostgreSQL instance. The following network types are supported:

- VPN
- Public network

Diagram

Figure 4-20 VPN network







Figure 4-21 Public network+SSL connection

Synchronization Process

Figure 4-22 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*.

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	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, and the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords)	The CONNECT permission for databases, the USAGE permission for schemas, the SELECT permission for tables, the SELECT permission for sequences, the SELECT permission for sequences, the SELECT permission for system table catalog pg_catalog.pg_authid (used for synchronizing user passwords), the UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create replication connections NOTE • The UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys are only used to temporarily lock tables to ensure data consistency after the migration. • To add the permission to create replication connections, perform the following steps: 1. Add host replication connections, perform the following steps: 1. Add host replication connections, perform the following steps: 2. Run select pg_reload_conf(); in the source database as user SUPERUSER, or restart the DB instance to apply the changes.

Synchronization Type	Full	Full+Incremental
Destination	 Database-level: The CREATEDB permission is required. Table-level: 	
	 To synchronize data permission is require 	abases, the CREATEDB red.
	and CREATE permis	thema, the CONNECT ssions for the database chema are required.
	contains the schem	on for the database that a, and the USAGE and s for the schema that
	 Synchronization user: permission is required 	
	privilege cannot be mo object permissions of	permissions: The default odified. Otherwise, the the destination database with those of the source

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 an RDS for PostgreSQL instance, the initial account can be used.

2. Network settings

Source database network settings:

You can synchronize data from on-premises PostgreSQL databases to RDS for 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.
- 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 4-23 Synchronization instance EIP

Synchronization instance created successfully. Its EIP is 10.
 Add this EIP to the destination database whitelist so that it can access the destination database.

IP of the synchronization instance is 192.
 Add this IP to the source database whitelist so that it can access the source database.

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.

Specinion in instance or Details ©
The stance of processes of the close of an end of the close of the clo

Figure 4-24 Synchronization instance information

Table 4-15 Task settings

Parameter	Description
Region	The region where the replication instance is deployed. You can change the region. To reduce latency and improve access speed, select the region closest to your workloads.
Project	The project corresponds to the current region and can be changed.
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 RDS for PostgreSQL instance you created.

Parameter	Description	
Synchronization Mode	 Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistence between the source and destination databases. 	
	 Full All objects and data in non-system databases are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable. 	
	– Incremental	
	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 remains 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 4-25 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

Parameter	Description
Database Username	A username for the source database.
Database Password	The password for the database username.
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 for 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 RDS for PostgreSQL instance.
Database Password	The password for the database username.

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

Figure 4-26 Synchronization mode



Table 4-19 Synchronization object

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

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.		
Synchronizati on 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 objects are selected before synchronization or select database-level synchronization.		


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.

Figure 4-27 Task Check



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

Figure 4-28 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 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 On-premises MySQL to GaussDB Distributed

Description


You can use real-time synchronization of DRS to synchronize on-premises MySQL to Huawei Cloud GaussDB. Full+incremental synchronization can ensure that data is always in sync between the source MySQL and the destination GaussDB.

Problems

- Enterprise workloads have been growing and evolving fast, and traditional databases lack the scalability needed to keep up. Enterprises need distributed databases.
- Building a traditional database means purchasing and installing servers, systems, databases, and other software. The O&M is expensive and difficult.

- Traditional databases are poor in complex queries.
- It is hard for traditional databases to smoothly synchronize data with no downtime.


Service architecture



Synchronization Principles

A full+incremental synchronization task includes the following operations:

- 1. In the full synchronization phase, tables, primary keys, and unique keys are synchronized.
- 2. Incremental data extraction is started to ensure that the incremental data generated during full data synchronization is completely extracted to the DRS instance.
- 3. A full synchronization is started.
- 4. An incremental synchronization is automatically started after the full synchronization is complete. The replay starts from the position where the full synchronization starts.
- 5. A comparison task is started after the incremental replay is complete to check the data consistency. Real-time comparison is supported.
- 6. Workloads synchronization is started if the data is consistent between the source and destination databases.



Service List

- Virtual Private Cloud (VPC)
- GaussDB
- Data Replication Service (DRS)
- Data Admin Service (DAS)

Notes on Usage

- The resource planning in this best practice is for demonstration only. Adjust it as needed.
- The end-to-end test data in this document is for reference only.
- Full synchronization is used to synchronize data. Incremental synchronization is used to synchronize data between the source and destination databases in real time.

Prerequisites

- You have registered with Huawei Cloud and completed account authentication.
- Your account balance is greater than or equal to \$0 USD.
- You have set up an on-premises MySQL database for testing.
- You have obtained the IP address, port number, account, and password of the MySQL database to be synchronized.

Resource Planning

Table 4-20 Resource planning

Categor y	Subcat egory	Planned Value	Remarks
VPC	VPC name	vpc-src-172	Specify a name that is easy to identify.
	Region	Test region	For low network latency and quick resource access, select the region nearest to you.
	AZ	AZ 3	-
	Subnet CIDR block	172.16.0.0/16	Select a subnet with sufficient network resources.
	Subnet name	subnet-src-172	Specify a name that is easy to identify.
On- premises MySQL	Databa se version	5.7.38	-
(source databas e)	Databa se user	test_info	Specify a database user. The following minimum permissions are required: SELECT, LOCK TABLES, REPLICATION SLAVE and REPLICATION CLIENT.
GaussDB	Instanc e name	Auto-drs-gaussdbv5- tar-1	Specify a name that is easy to identify.
	Databa se version	GaussDB 1.3 Enterprise Edition	-
	Instanc e type	Distributed (1 CN, 3 DN shards, and 3 replicas)	Select a distributed instance for the test.
	Deploy ment model	Independent	-
	Transa ction consist ency	Strong consistency	-
	Shards	3	-

Categor y	Subcat egory	Planned Value	Remarks
	Coordi nator nodes	3	-
	Storag e type	Ultra-high I/O	-
	AZ	AZ 2	In this example, a single AZ is select. You are advised to select multiple AZs to improve instance availability in actual use.
	Instanc e specific ations	General-enhanced II 8 vCPUs 64 GB	Small specifications are selected for this test instance. You are advised to configure specifications based on service requirements in actual use.
	Storag e space	480 GB	A small storage space is selected for this test instance. You are advised to configure the storage space based on service requirements in actual use.
	Disk encryp tion	Disable	In this example, disk encryption is disabled. Enabling disk encryption improves the security of data, but may slightly affect the database read/write performance.
Logging in to the	DB engine	GaussDB	-
databas e through DAS	Databa se source	GaussDB	Select the GaussDB instance created in this example.
	Databa se name	postgres	-
	Userna me	root	-
	Passwo rd	-	Password of the root user of the GaussDB instance created in this example

Categor y	Subcat egory	Planned Value	Remarks
DRS synchron ization task	Task name	DRS-test-info	Specify a name that is easy to identify.
	Destin ation databa se name	test_database_info	Specify a name that is easy to identify. The name must be compatible with the MySQL database name.
	Source DB engine	MySQL	-
	Destin ation DB engine	GaussDB	-
	Netwo rk type	Public network	Public network is used in this example.

Flowchart

Figure 4-29 shows the main operation flowchart.

Create a VPC and security group.

Create a distributed GaussDB instance.

Construct test data.

Create a DRS synchronization task.


Figure 4-29 Flowchart

Creating a VPC


Create a VPC to prepare network resources for creating a GaussDB instance.

- Step 1 Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.



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


Create a security group for creating a GaussDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click oin the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.


The VPC console is displayed.



- **Step 4** Choose **Access Control** > **Security Groups**.
- **Step 5** Click **Create Security Group**.
- **Step 6** Specify a security group name and other information.



- Step 7 Click OK.
- **Step 8** Return to the security group list and click the security group name (**sg-01** in this example).
- Step 9 Click the Inbound Rules tab and then click Add Rule.
- **Step 10** Configure an inbound rule, add the IP address of the source database, and click **OK**.




----End


Creating a Distributed GaussDB Instance

This section describes how to create a distributed GaussDB instance as the destination database.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Under the service list, choose **Databases** > **GaussDB**.
- Step 4 Click Buy DB Instance.
- **Step 5** Configure basic information for the instance.



Step 6 Configure instance specifications.




Select small specifications for this test instance. You are advised to configure specifications based on service requirements in actual use.

Step 7 Select a VPC and security group (created in Creating a VPC and Creating a Security Group) for the instance and configure the database port.



Step 8 Configure password and other information.



- **Step 9** Click **Next**, confirm the information, and click **Submit**.
- **Step 10** Go to the instance list.

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

----End

Constructing Test Data

Before the synchronization, prepare some data types in the source database for verification after the synchronization is complete.

For details about the data types supported by DRS, see MySQL->GaussDB.

Perform the following steps to construct data in the source database:

- **Step 1** Use a database connection tool to connect to the source MySQL database based on its IP address.
- **Step 2** Construct data in the source database based on data types supported by DRS.
 - 1. Create a test user.

create user test info identified by xxx;

test_info indicates the user created for the test, and *xxx* indicates the password of the user.

2. Create a database named **test info** under the user.

CREATE DATABASE test_info;

3. Create a table in the **test_info** database.

CREATE TABLE `test_info`.`test_table` (

- 'id' int NOT NULL,
- `c1` char(10) DEFAULT NULL,
- `c2` varchar(10) DEFAULT NULL,
- `c3` binary(10) DEFAULT NULL,
- `c4` varbinary(10) DEFAULT NULL,
- `c5` tinyblob,
- `c6` mediumblob,
- `c7` longblob,
- `c8` tinytext,
- `c9` text,
- `c10` mediumtext,
- `c11` longtext,
- `c12` enum('1','2','3') DEFAULT NULL,
- `c13` set('1','2','3') DEFAULT NULL,
- `c14` tinyint DEFAULT NULL,
- `c15` smallint DEFAULT NULL,
- `c16` mediumint DEFAULT NULL,
- `c17` bigint DEFAULT NULL,
- `c18` float DEFAULT NULL,
- `c19` double DEFAULT NULL,
- `c20` date DEFAULT NULL,
- `c21` datetime DEFAULT NULL,
- `c22` timestamp,
- `c23` time DEFAULT NULL,
- `c24` year DEFAULT NULL,
- `c25` bit(10) DEFAULT NULL,
- `c26` json DEFAULT NULL,
- `c27` decimal(10,0) DEFAULT NULL,
- `c28` decimal(10,0) DEFAULT NULL,
- PRIMARY KEY ('id')

);

4. Assign permissions to the user.

GRANT SELECT, LOCK TABLES ON <database>. to test_info; GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* to test info;

In the preceding command, <code>test_info</code> indicates the user created for this test, <code><database></code> indicates the name of the database to be synchronized, and <code></code> indicates the name of the table to be synchronized. Replace them based on the site requirements.

5. Insert two rows of data into the table.

insert into test info.test table values

(1,'a','b','111','111','tinyblob','mediumblob','longblob','tinytext','text','mediumtext','longtext','1','3',1,2,3,4,1.123,1.1234,'2024-03-08','2024-03-08 08:00:00','2024-03-08

08:00:00','08:00:00','2024','1010','{"a":"b"}',1.23,1.234);


insert into test_info.test_table values

(2,'a','b','111','111','tinyblob','mediumblob','longblob','tinytext','text','med iumtext','longtext','1','3',1,2,3,4,1.123,1.1234,'2024-03-08','2024-03-08 08:00:00','2024-03-08

08:00:00','08:00:00','2024','1010','{"a":"b"}',1.23,1.234);

Step 3 Create a database in the destination GaussDB instance.

- 1. Log in to the management console.
- 2. Click in the upper left corner and select a region.
- 3. Click in the upper left corner of the page and choose **Databases** > **Data Admin Service**.
- 4. In the navigation pane on the left, click **Development Tool** to go to the login list page.
- 5. Click Add Login.
- 6. On the displayed page, select the DB engine, source database, and target DB instance, enter the login username, password, and description (optional), and enable Collect Metadata Periodically and Show Executed SQL Statements.
 If Collect Metadata Periodically is enabled, select Remember Password.
- 7. Click **Test Connection** to check whether the connection is successful. If a message is displayed indicating connection successful, continue with the operation. If a message is displayed indicating connection failed and the failure cause is provided, make modifications according to the error message.
- 8 Click OK
- 9. Locate the added instance, click **Log In** in the **Operation** column.



Data
SOL Operations
Database Management Import and Export Account Management

Home SQL History X SQL Window X

10. Choose **SQL Operations** > **SQL Window** on the top menu bar.

11. Run the following statement to create a database compatible with MySQL: test_database_info indicates the database name. Replace it based on the site requirements.
CREATE DATABASE test_database_info DBCOMPATIBILITY 'mysql';

----End

Performing a Pre-Check

Before creating a task, check whether synchronization conditions are met.


Before synchronization, refer to **Precautions**.

Creating a DRS Synchronization Task

This section describes how to create a DRS instance and synchronize data from the **test_info** database in the on-premises MySQL database to the **test_database_info** database in the GaussDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.

 Select the region where the destination instance is deployed.
- **Step 3** Click the service list icon on the left and choose **Databases** > **Data Replication Service**.
- **Step 4** In the navigation pane on the left, choose **Data Synchronization Management**. On the displayed page, click **Create Synchronization Task**.
- **Step 5** Configure synchronization instance information.
 - 1. Select a region, and project, and enter a task name.



2. Specify Data Flow, Source DB Engine, Destination DB Engine, Network Type, DRS Task Type, Destination DB Instance, Synchronization Instance

Subnet (optional), **Synchronization Mode**, **Specifications**, **AZ**, and **Tags** (optional).



- 3. Click Create Now.
- **Step 6** Configure the source and destination database information.
 - 1. Enter the IP address, port number, username, and password of the source database.

Click **Test Connection**.




2. Enter the username and password of the destination database.

Click Test Connection.



3. Click Next. In the displayed box, read the message carefully and click Agree.



Step 7 Configure the synchronization task.

1. Select the object type for full synchronization. If the table structure to be synchronized has not been created in the destination database, select **Table**


structure (the table structure contains primary keys and unique keys) for **Synchronization Object Type**. Otherwise, deselect **Table structure**. Select **Index** for **Synchronization Object Type** based on the site requirements.




- 2. Specify Incremental Conflict Policy.
 - Ignore: The system will ignore the conflicting data and continue the subsequent synchronization process. If you select Ignore, data in the source database may be inconsistent with that in the destination database.
 - **Report error**: The synchronization task will be stopped and fail. You can view the details in synchronization logs.
 - Overwrite: Conflicting data will be overwritten.



3. Select the databases and tables of the source database to be synchronized. In this test, select the **test_table** table from the **test_info** database.




4. Locate the database and table, respectively, and click **Edit** to change the database name and table name.




5. On the displayed box, enter a new name, for example, **DATATYPELIST_After**.

The name cannot include special characters. Otherwise, an error will be reported during SQL statement execution after the synchronization.




6. Confirm the settings and click **Next**.



Step 8 Confirm advanced settings.


The information on the **Advanced Settings** page is for confirmation only and cannot be modified. After confirming the information, click **Next**.




Step 9 Process data.

On the **Processing Columns** tab, select the column to be synchronized and change its name, for example, change **c1** to **new-line**.

1. Click **Edit** next to the table to be processed.




2. Edit the c1 column.



- 3. Enter the new name **new-line** and click **Confirm**.
- 4. Click Next.

Step 10 Perform a pre-check.

- 1. After all settings are complete, perform a pre-check to ensure that the synchronization is successful.
- 2. If any check item fails, review the cause and rectify the fault. Then, click **Check Again**.




3. If all check items pass the pre-check, click **Next**.

Step 11 Confirm the task.

1. Check that all configured information is correct.



- 2. Click **Submit**. In the display box, select **I have read the precautions**.
- 3. Click Submit.



Step 12 After the task is submitted, view and manage it.

After the task is created, return to the task list to view the status of the created task.

----End

Verifying Data After Synchronization

When the task status changes to **Incremental**, the full synchronization is complete. You can log in GaussDB and view the data synchronization result.


Step 1 Wait until the synchronization task status becomes **Incremental**.



- **Step 2** Click the task name to go to the **Basic Information** page.
- **Step 3** Verify data consistency.
 - 1. Choose **Synchronization Comparison** > **Object-Level Comparison** to view the database and table synchronization results.



 Choose Synchronization Comparison > Data-Level Comparison, click Create Comparison Task, and view the synchronization results of the rows in the table.



Step 4 Connect to test_database_info in GaussDB using DAS.

For details about how to connect to an instance through DAS, see **Adding Login Information**.

Step 5 Run the following statement to query the full synchronization result:

SELECT * FROM test info.datatypelist after;

After the schema in MySQL is synchronized, it will be used as the schema in GaussDB. Therefore, it is required to add the schema in the query statement for exact query.

All data types in the table were successfully synchronized and the data is correct.

Step 6 Verify incremental synchronization.

In full+incremental synchronization, after the full synchronization is complete, the data that is written to the source database after the task is created can still be synchronized to the destination database until the task is stopped. The following describes how to synchronize incremental data from the source database to the destination database:

- 1. Use a database connection tool to connect to the source MySQL database based on its IP address.
- 2. Run the following statement to insert a data record into the source database: Insert a data record whose ID is 3.

```
insert into test_info.test_table values (3,'a','b','111','111','tinyblob','mediumblob','longblob','tinytext','text','mediumtext','longtext','1','3',1,2,3,4, 1.123,1.1234,'2024-03-08','2024-03-08 08:00:00','2024-03-08 08:00:00','08:00:00','2024','1010','{"a":"b"}',1.23,1.234);
```

3. Run the following statement in the destination database to query the result: SELECT * FROM test_info.datatypelist_after;

The new data in the source database has been synchronized to the destination database in real time.

Step 7 Stop the synchronization task.

After data is completely synchronized to the destination database, stop the synchronization task.

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

2. In the display box, click Yes.

----End

4.5 From On-premises Oracle to GaussDB Distributed

4.5.1 Purpose


Description

You can use real-time synchronization of DRS to synchronize on-premises Oracle to Huawei Cloud GaussDB. Full+incremental synchronization can ensure that data is always in sync between the source Oracle and the destination GaussDB.

Problems

- With the rapid increase of enterprise workloads, traditional databases have poor scalability and distributed databases are required.
- Building traditional databases require purchasing and install servers, systems, databases, and other software. Its O&M is expensive and difficult.
- The performance of complex queries for traditional databases is poor.
- It is hard for traditional databases to smoothly migrate data without interrupting services.

Migration Architecture




Migration Principles

Perform the following operations to complete full and incremental synchronization:

1. In the full synchronization phase, migrate tables, primary keys, and unique keys.

- 2. Start incremental data extraction to ensure that the incremental data generated during full data synchronization is completely extracted to the DRS instance.
- 3. Start the full migration task.
- 4. Automatically perform incremental synchronization after the full migration is complete. The playback starts from the location where the full migration starts.
- 5. Start the comparison task after the incremental replay is complete to check the data consistency. Real-time comparison is supported.
- 6. Start migration if the data is consistent.



Service List

- Virtual Private Cloud (VPC)
- GaussDB
- Data Replication Service (DRS)
- Data Admin Service (DAS)

Notes on Usage

- The resource planning in this document is for demonstration only. Adjust it as needed.
- The end-to-end test data in this document is for reference only.
- Full synchronization is used to migrate data. Incremental synchronization is used to synchronize data between the source and destination databases in real time.

Prerequisites

- You have registered with Huawei Cloud and completed account authentication.
- Your account balance is greater than or equal to \$0 USD.
- You have set up an on-premises Oracle database for testing.
- You have obtained the IP address, port number, account, and password of the Oracle database to be migrated.

4.5.2 Resource Planning

Table 4-21 Resource planning

Categor y	Subcat egory	Planned Value	Remarks
VPC	VPC name	vpc-src-172	Customize a name for easy identification.
	Region	Test region	For low network latency and quick resource access, select the region nearest to you.
	AZ	AZ3	-
	Subnet	172.16.0.0/16	Select a subnet with sufficient network resources.
	Subnet name	subnet-src-172	Customize a name for easy identification.
Oracle	Name	orcl	Customize a name for easy identification.
	Specifi cations	16 vCPUs 32 GB	-
	Databa se version	11.2.0.1	-
	Databa se user	test_info	Customize a user. However, the user must have the following permissions during migration: CREATE SESSION, SELECT ANY TRANSACTION, SELECT ANY TABLE, SELECT ANY DICTIONARY, and EXECUTE_CATALOG_ROLE.
GaussDB	Instanc e name	Auto-drs-gaussdbv5- tar-1	Customize a name for easy identification.
	Databa se version	GaussDB 1.3 Enterprise Edition	-
	Instanc e type	Distributed (1 CN, 3 DN shards, and 3 replicas)	Select a distributed instance for the test.
	Deploy ment model	Independent	-

Categor y	Subcat egory	Planned Value	Remarks
	Transa ction consist ency	Strong consistency	-
	Shards	3	-
	Coordi nator nodes	3	-
	Storag e type	Ultra-high I/O	-
	AZ	AZ2	Select a single AZ for the test. You are advised to select multiple AZs to improve instance availability in actual use.
	Instanc e specific ations	General-enhanced II 8 vCPUs 64 GB	Select small specifications for the test. You are advised to configure specifications based on service requirements in actual use.
	Storag e space	480G	Select a small storage space for the test. You are advised to configure storage space based on service requirements in actual use.
	Disk encryp tion	Disable	In this example, disk encryption is disabled. Enabling disk encryption improves data security, but slightly affects the read and write performance of the database.
Logging in to the	DB engine	GaussDB	-
databas e through DAS	Databa se source	GaussDB	Select the GaussDB instance created in this example.
	Databa se name	postgres	-
	Userna me	root	-

Categor y	Subcat egory	Planned Value	Remarks	
	Passwo rd	-	Password of the root user of the GaussDB instance created in this example	
DRS migratio n task	Migrati on task name	DRS-test-info	Customize a name for easy identification.	
	Destin ation databa se name	test_database_info	Customize a name for easy identification, but the name must be compatible with the Oracle database name.	
	Source DB engine	Oracle	-	
Destin ation DB engine		GaussDB	-	
	Netwo rk type	Public	Select the public network for the test.	

4.5.3 Operation Process

Figure 4-30 shows the main operation flowchart.




Figure 4-30 Operation flowchart


4.5.4 Creating a VPC and Security Group

Create a VPC and security group for a GaussDB instance.


Creating a VPC

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.



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 oin the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.




- **Step 4** 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-01** in this example).
- **Step 9** Click the **Inbound Rules** tab, and then click **Add Rule**.



Step 10 Configure an inbound rule, add the IP address of the source database, and click **OK**.




----End


4.5.5 Creating a GaussDB Instance

This section describes how to create a GaussDB instance as the destination database.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Under the service list, choose **Databases** > **GaussDB**.
- Step 4 Click Buy DB Instance.
- **Step 5** Configure the instance name and basic information.



Step 6 Configure instance specifications.




Select small specifications for the test. You are advised to configure specifications based on service requirements in actual use.

Step 7 Select a VPC and security group (created in **Creating a VPC and Security Group**) for the instance and configure the database port.



Step 8 Configure password and other information.



- Step 9 Click Next, confirm the information, and click Submit.
- Step 10 Go to the instance list.

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

----End

4.5.6 Generating Test Data

Before the synchronization, prepare some data types in the source database for verification after the synchronization is complete.

The following table lists data types supported by DRS.

Table 4-22 Data type mapping

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
CHAR	character	Suppor ted	Supporte d	Supported . The spaces before and after the character are ignored.	Supported. The spaces before and after the character are ignored.	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
VARCH AR	character varying	Suppor ted	Supporte d	Supported	Supported	The precision ranges of the source and destination databases are different, causing precision loss.
VARCH AR2	character varying	Suppor ted	Supporte d	Supported	Supported	-
NCHAR	character	Suppor ted	Supporte d	Supported . The spaces before and after the character are ignored.	Supported. The spaces before and after the character are ignored.	-
NVARC HAR2	nvarchar2	Suppor ted	Supporte d	Supported	Supported	-
NUMBE R	numeric	Suppor ted	Supporte d	Supported	Supported	-
NUMBE R (6,3)	numeric(6,3)	Suppor ted	Supporte d	Supported	Supported	-
NUMBE R (6,0)	Integer	Suppor ted	Supporte d	Supported	Supported	-
NUMBE R (3)	smallint	Suppor ted	Supporte d	Supported	Supported	-
NUMBE R (6,-2)	integer	Suppor ted	Supporte d	Supported	Supported	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
BINARY _FLOAT	real	Unsup ported (The destin ation databa se does not suppor t creatin g tables using the primar y key.)	Supporte	Unsupport	Supported	The precision ranges of the source and destination databases are different, causing precision loss.
BINARY _DOUB LE	double precision	Unsup ported (The destin ation databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport ed	Supported	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
FLOAT	real	Unsup ported (The destin ation databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport ed	Supported	The precision ranges of the source and destination databases are different, causing precision loss.
INT	numeric	Suppor ted	Supporte d	Supported	Supported	-
INTEGE R	numeric	Suppor ted	Supporte d	Supported	Supported	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
DATE	date	Supported	Supporte	Unsupport	Supported	If a table with date type is created in the destination database, the data type precision range in the source database is different from that in the destination database, causing precision loss. Therefore, compariso n is not supported.
TIMEST AMP	timestam p(6) without time zone	Suppor ted	Supporte d	Unsupport ed	The value is accurate to six decimal places.	The maximum precision supported by the source database is 6.

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
TIMEST AMP_T Z	timestam p(6) with time zone	Unsup ported (The source databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport ed	Filter this column.	-
TIMEST AMP_LT Z	timestam p(6) with time zone	Unsup ported (The destin ation databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport	Filter this column.	-
INTERV AL_YM	interval year to month	Suppor ted	Supporte d	Unsupport ed	Unsupport ed	Incrementa l synchroniz ation does not support this type.

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
INTERV AL_DS	interval day to second	Suppor ted	Supporte d	Unsupport ed	Unsupport ed	Incrementa l synchroniz ation does not support this type. The maximum precision supported by the source database is 6.
BLOB	bytea	Unsup ported (The source databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport ed	Filter this column.	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
CLOB	text	Unsup ported (The source databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport	Filter this column.	
NCLOB	text	Unsup ported (The source databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport ed	Filter this column.	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
LONG	text	Unsup ported (The source databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport	Filter this column.	-
LONG_ RAW	bytea	Unsup ported (The source databa se does not suppor t creatin g tables using the primar y key.)	Supporte d	Unsupport ed	Filter this column.	-

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
RAW	bytea	Unsup ported (The destin ation databa se does not suppor t creatin g tables using the primar y key.)	Supporte	Unsupport ed	Supported	-
RowID	character varying(1 8)	Suppor ted	Supporte d	Supported	Supported	-
BFILE	-	Unsup ported	Unsuppo rted	Unsupport ed	Unsupport ed	Restriction s on the source database: The bfile type is not supported.
XMLTYP E	-	Unsup ported	Unsuppo rted	Unsupport ed	Unsupport ed	Restriction s on the source database: The xmltype type is not supported.

Source Data Type	Destinati on Data Type	Sync (Sourc e Data Type as Primar y Key)	Sync (Source Data Type as Non- Primary Key)	Comparis on (Source Data Type as Primary Key)	Comparis on (Source Data Type as Non- Primary Key)	Remarks
UROWI D	-	Unsup ported	Unsuppo rted	Unsupport ed	Unsupport ed	Full and incrementa l synchroniz ations are not supported.
sdo_ge ometry	-	Unsup ported	Unsuppo rted	Unsupport ed	Unsupport ed	Restriction s on the source database: The sdo_geome try type is not supported.
NUMBE R(*, 0)	numeric	Suppor ted	Supporte d	Supported	Supported	-

Perform the following steps to generate data in the source database:

- **Step 1** Use a database connection tool to connect to the source Oracle database based on its IP address.
- **Step 2** Construct data in the source database based on data types supported by DRS.
 - 1. Create a test user.
 - create user test_info identified by xxx;

test_info indicates the user created for the test, and *xxx* indicates the password of the user.

2. Assign permissions to the user.

grant dba to test_info;

3. Create a data table under the user.

CREATE TABLE test_info.DATATYPELIST(
ID INT,
COL_01_CHAR_____E CHAR(100),

COL_02_NCHAR____E NCHAR(100),

COL_03_VARCHAR__E VARCHAR(1000),

COL_04_VARCHAR2_E VARCHAR2(1000),

```
COL_05_NVARCHAR2_E NVARCHAR2(1000),
COL_06_NUMBER___E NUMBER(38,0),
COL_07_FLOAT____E FLOAT(126),
COL 08 BFLOAT E BINARY FLOAT,
COL_09_BDOUBLE___E BINARY_DOUBLE,
COL_10_DATE_____E DATE DEFAULT SYSTIMESTAMP,
COL_11_TS____E TIMESTAMP(6),
COL_12_TSTZ_____E TIMESTAMP(6) WITH TIME ZONE,
COL 13 TSLTZ E TIMESTAMP(6) WITH LOCAL TIME ZONE,
COL_14_CLOB_____E CLOB DEFAULT EMPTY_CLOB(),
COL 15 BLOB E BLOB DEFAULT EMPTY BLOB(),
COL_16_NCLOB____E NCLOB DEFAULT EMPTY_CLOB(),
COL_17_RAW_____E RAW(1000),
COL_19_LONGRAW___E LONG RAW,
COL 24 ROWID E ROWID,
PRIMARY KEY(ID)
):
```

4. Insert two rows of data.

insert into test_info.DATATYPELIST values(4,'huawei','xian','Shannxi','zhongguo','shijie', 666,12.321,1.123,2.123,sysdate,sysdate,sysdate,sysdate,'hw','cb','df','FF','FF ','AAAYEVAAJAAAACrAAA');

insert into test_info.DATATYPELIST values(2,'Migrate-test','test1','test2','test3','test4', 666,12.321,1.123,2.123,sysdate,sysdate,sysdate,sysdate,'hw','cb','df','FF','FF','AAAYEVAAJAAAACrAAA');

5. Make the above statements take effect.

commit;

Step 3 Create a database at the destination server.

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select a region.
- 3. Click in the upper left corner of the page and choose **Databases** > **Data Admin Service**.
- 4. In the navigation pane on the left, click **Development Tool** to go to the login list page.
- 5. Click Add Login.
- 6. On the displayed page, select the DB engine, source database, and target DB instance, enter the login username, password, and description (optional), and enable Collect Metadata Periodically and Show Executed SQL Statements.


 If Collect Metadata Periodically is enabled, select Remember Password.
- 7. Click **Test Connection** to check whether the connection is successful.

If a message is displayed indicating connection successful, continue with the operation. If a message is displayed indicating connection failed and the failure cause is provided, make modifications according to the error message.

- 8. Click OK.
- 9. Locate the added instance, click **Log In** in the **Operation** column.



10. Choose **SQL Operations** > **SQL Window** on the top menu bar.



11. Run the following statement to create a database compatible with Oracle: **test_database_info** indicates the database name. Replace it based on the site requirements.

CREATE DATABASE test database info DBCOMPATIBILITY 'ORA';

----End

4.5.7 Synchronizing Databases

This section describes how to create a DRS instance and synchronize data from the **test_info** database on the local Oracle to the **test_database_info** database on GaussDB instance.

Pre-Check

Before creating a task, check the synchronization conditions.

Before synchronization, refer to **Precautions**.

Creating a Synchronization Task

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.

 Select the region where the destination instance is deployed.
- **Step 3** Click the service list icon on the left and choose **Databases** > **Data Replication Service**.

- **Step 4** In the navigation pane on the left, choose **Data Synchronization Management**. On the displayed page, click **Create Synchronization Task**.
- **Step 5** Configure synchronization instance information.
 - 1. Select a region, and project, and enter a task name.




 Specify Data Flow, Source DB Engine, Destination DB Engine, Network Type, DRS Task Type, Destination DB Instance, Synchronization Instance Subnet (optional), Synchronization Mode, Specifications, AZ, and Tags (optional).

Figure 4-31 Synchronization instance details



- 3. Click Create Now.
- **Step 6** Configure the source and destination database information.
 - 1. Enter the IP address, port number, username, and password of the source database.


Click Test Connection.



Enter the username and password of the destination database. Click **Test Connection**.




3. Click **Next**. In the displayed box, read the message carefully and click **Agree**.




Step 7 Configure the synchronization task.


1. Select the databases and tables of the source database to be migrated. For example, select the **DATATYPELIST** table from the **test_info** database.




2. Locate the database and table, respectively, and click **Edit** to change the database name and table name.



On the displayed box, enter a new name, for example, DATATYPELIST_After.
 The name cannot include special characters. Otherwise, an error will be reported during SQL statement execution after the migration.




4. Confirm the settings and click Next.



Step 8 Confirm advanced settings.


The information on the **Advanced Settings** page is for confirmation only and cannot be modified. After confirming the information, click **Next**.




Step 9 Process data.

On this page, you can process the table to be migrated. Select the column to be migrated and change its name, for example, change **COL_01_CHAR____E** to **new-line**.


1. Select the table to be processed.



2. Edit the COL_01_CHAR____E column.




3. Enter the new name **new-line** and click **Confirm**.




4. Click Next.

Step 10 Perform a pre-check.

1. After all settings are complete, perform a pre-check to ensure that the migration is successful.




2. If any check fails, review the cause and rectify the fault. Then, click **Check Again**.




3. If all check items pass the pre-check, click **Next**.

Step 11 Confirm the task.

1. Check whether all configured information is correct.



- 2. Click **Start Task**. In the display box, select **I have read the precautions**.
- 3. Click Submit.



Step 12 After the task is submitted, view and manage it.

After the task is created, return to the task list to view the status of the created task.

----End

4.5.8 Verifying Data


When the task status changes to **Incremental**, the full synchronization is complete. You can log in GaussDB and view the data migration result.

Step 1 Wait until the migration task status becomes **Incremental**.



- **Step 2** Click the task name to go to the **Basic Information** page.
- **Step 3** On the **Synchronization Progress** page, view the full synchronization result.

As shown in the following figure, the **DATATYPELIST** table in the **TEST_INFO** database has been migrated to **shard_0**. Two rows were migrated successfully.




Step 4 Verify data consistency.

1. Choose **Synchronization Comparison** > **Object-Level Comparison** to view the database and table migration results.



2. Choose Synchronization Comparison > Data-Level Comparison, click Create Comparison Task, and view the migration results of the rows in the table.



Step 5 Connect to the destination database **test_database_info** through DAS.


For details about how to connect to an instance through DAS, see **Adding Login Information**.

Step 6 Query the full synchronization result.

SELECT * FROM test_info.datatypelist_after;

After the schema in Oracle is migrated, it will be used as the schema in GaussDB. Therefore, it is required to add the schema in the query statement for exact query.

As shown in the following figure, all data types in the table were successfully migrated and the data is correct.



Step 7 Verify incremental synchronization.

In the migration type of full and incremental synchronization, after the full synchronization is complete, the data that is written to the source database after the task is created can still be synchronized to the destination database until the task is complete. The following describes how to synchronize incremental data from the source database to the destination database.

- 1. Use a database connection tool to connect to the source Oracle database based on its IP address.
- 2. Run the following statement to insert a data record into the source database: Insert a data record whose ID is 1.

```
insert into test_info.DATATYPELIST values(1,'Migrate-test','test1','test2','test3','test4', 666,12.321,1.123,2.123,sysdate,sysdate,sysdate,sysdate,'hw','cb','df','FF','AAAYEVAAJAAAACrAAA'); commit:
```

3. Run the following statement in the destination database to query the result: SELECT * FROM test_info.datatypelist_after;

As shown in the following figure, the new data inserted in the source database has been synchronized to the destination database in real time.




Step 8 Stop the migration task.

After data is completely migrated to the destination database, stop the migration task.

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



2. In the display box, click Yes.



----End

4.6 From On-premises Oracle to GaussDB Centralized

4.6.1 Purpose

Description

This practice uses UGO syntax migration and DRS real-time synchronization to synchronize on-premises Oracle to Huawei Cloud GaussDB. You can create a full +incremental synchronization task to synchronize data from the source Oracle to the destination GaussDB.

Problems


- With the rapid increase of enterprise workloads, traditional databases have poor scalability and distributed databases are required.
- Building traditional databases require purchasing and install servers, systems, databases, and other software. Its O&M is expensive and difficult.
- The performance of complex queries for traditional databases is poor.

 It is hard for traditional databases to smoothly synchronize data with no downtime

Migration Principles

In this practice, UGO is used to migrate heterogeneous database schemas, and DRS is used to only synchronize data. The principles are as follows:

- 1. UGO is used to synchronize basic objects.
- 2. DRS is used to perform full synchronization and suspend the task after the task enters the incremental phase.
- 3. UGO is used to synchronize common indexes.
- 4. DRS is used to resume incremental synchronization. After data is synchronized, stop services on the source end for data comparison to ensure data consistency. Then, you need to stop the DRS task.
- 5. Finally, UGO is used to migrate triggers, events, jobs, foreign keys, and sequences, and services are migrated to GaussDB.



Service List

- Virtual Private Cloud (VPC)
- GaussDB
- Database and Application Migration UGO (UGO)
- Data Replication Service (DRS)
- Data Admin Service (DAS)

Notes on Usage

- The resource planning in this document is for demonstration only. Adjust it as needed.
- The end-to-end test data in this document is for reference only.
- Full synchronization is used to migrate data. Incremental synchronization is used to synchronize data between the source and destination databases in real time.

Prerequisites

- You have registered with Huawei Cloud and completed account authentication.
- Your account balance is greater than or equal to \$0 USD.
- You have set up an on-premises Oracle database for testing.
- You have obtained the IP address, port number, account, and password of the Oracle database to be migrated.

4.6.2 Resource Planning

Table 4-23 Resource planning

Categor y	Subcat egory	Planned Value	Remarks
VPC	VPC name	vpc-src-172	Customize a name for easy identification.
	Region	Test region	For low network latency and quick resource access, select the region nearest to you.
	AZ	AZ3	-
	Subnet	172.16.0.0/16	Select a subnet with sufficient network resources.
	Subnet name	subnet-src-172	Customize a name for easy identification.
Oracle	Name	orcl	Customize a name for easy identification.
	Specifi cations	16 vCPUs 32 GB	-
	Databa se version	11.2.0.1	-
	Databa se user	test_info	Customize a user. However, the user must have the following permissions during migration: CREATE SESSION, SELECT ANY TRANSACTION, SELECT ANY TABLE, SELECT ANY DICTIONARY, and EXECUTE_CATALOG_ROLE.
GaussDB	Instanc e name	gauss-b193-cent	Customize a name for easy identification.

Categor y	Subcat egory	Planned Value	Remarks
	Databa se version	GaussDB 2.7 Enterprise Edition	-
	Instanc e type	Primary/Standby	In this example, a primary/ standby instance is used.
	Deploy ment model	Independent	-
	Transa ction consist ency	Strong consistency	-
	Storag e type	Ultra-high I/O	-
	AZ	AZ2	Select a single AZ for the test. You are advised to select multiple AZs to improve instance availability in actual use.
	Instanc e specific ations	General-enhanced II 8 vCPUs 64 GB	Select small specifications for the test. You are advised to configure specifications based on service requirements in actual use.
	Storag e space	480G	Select a small storage space for the test. You are advised to configure storage space based on service requirements in actual use.
	Disk encryp tion	Disable	In this example, disk encryption is disabled. Enabling disk encryption improves data security, but slightly affects the read and write performance of the database.
Logging in to the databas e through DAS	DB engine	GaussDB	-
	Databa se source	GaussDB	Select the GaussDB instance created in this example.


Categor y	Subcat egory	Planned Value	Remarks
	Databa se name	postgres	-
	Userna me	root	-
	Passwo rd	-	Password of the root user of the GaussDB instance created in this example
UGO migratio n task	Databa se evaluat ion project	Oracle-Centralized	Customize a name for easy identification.
	Object migrati on project	Oracle-GaussDB- Centralized	Customize a name for easy identification.
	Source DB engine	Oracle 11g	-
	Target DB engine	GaussDB Primary/ Standby - 2.8	
	Netwo rk type	Public Network	Select the public network for the test.
DRS migratio n task	Migrati on task name	DRS-test-info	Customize a name for easy identification.
	Destin ation databa se name	test_database_info	Customize a name for easy identification, but the name must be compatible with the Oracle database name.
	Source DB engine	Oracle	-
	Destin ation DB engine	GaussDB Centralized	-

Categor y	Subcat egory	Planned Value	Remarks
	Netwo rk type	Public	Select the public network for the test.

4.6.3 Operation Process

Figure 4-32 shows the main operation flowchart.

Figure 4-32 Flowchart



4.6.4 Preparations

Permissions Required for UGO

• The permission check for the source database must be passed. Table 4-24 lists the permissions that need to be checked when the source database type is Oracle.

Table 4-24 Required checks

Checked Item	Description	Mandatory
DBMS_METADATA	Permission to retrieve metadata from the Oracle database dictionary. This permission is used to obtain the DDL of schema objects.	Yes
Dynamic View	Permission to access various dynamic performance views. This permission is used to obtain basic database information.	Yes
Schema Objects	Permission to check schema objects to be evaluated. At least one object needs to be evaluated.	Yes
DBA	DBA permission required for subsequent operations.	No If Check Result is Alarm, some objects could not be collected because of permissions, but the evaluation project can still be created successfully.

- When connecting to the target database, you must have the permissions needed to create, delete, and modify databases objects, such as schemas, tables, programs, indexes, users, functions, and views.
- You have the permissions needed to create an evaluation project. For details, see Permission Management.

For more information about permissions required for UGO, see Here.

Permissions Required for DRS

The permissions required by DRS for connecting to the Oracle database must meet the following requirements to ensure smooth data synchronization:

- Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser).
- Full+incremental synchronization and incremental synchronization:
 - Oracle 12c or later in tenant mode:

To synchronize a container database (CDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.

To synchronize a pluggable database (PDB) of Oracle 12c or later, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, and CREATE SESSION, SELECT ANY DICTIONARY, EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL) permissions for a CDB.

- Oracle 12c or later in non-tenant mode:
 - You must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, and LOGMINING.
- To synchronize a database of Oracle 11g or earlier, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT ON <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, and SELECT ANY TRANSACTION.
- During incremental synchronization, enable PK, UI, or ALL supplemental logging for the source Oracle database at the database level or table level. If supplemental logging is enabled at table level, enable supplemental logging again after you rebuild or rename tables. During the synchronization, ensure that the preceding settings are always enabled.
- Oracle 12c or later does not support incremental synchronization using accounts whose ORACLE_MAINTAINED is Y (except system/sys), because accounts with this attribute do not have the permission to parse logs.
- The permissions required by DRS for connecting to the GaussDB database must meet the following requirements to ensure smooth data synchronization:
 - Database-level permissions: Log in to the postgres base database as user root or user DATABASE with the Sysadmin role, and grant the CREATE and CONNECT permissions to user DATABASE.
 - Statement: **GRANT CREATE, CONNECT ON DATABASE <database> TO <user>**

- Schema-level permissions: Log in to the database as user root or user DATABASE with the Sysadmin role, or the owner of the database, and grant the CREATE and USAGE permissions of the schema to the user.
 - Statement: GRANT CREATE, USAGE ON SCHEMA <schema> TO <user>
- Table-level permissions: Log in to the database as user **root** or user **DATABASE** with the Sysadmin role, or the owner of the database, and grant the DML permission for tables in the schema to the user. (The SELECT permission is required only for tables without primary keys.)

Statement for granting the DML permission on all tables in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON ALL TABLES IN SCHEMA <schema> TO <user>

Statement for granting the DML permission on a specified table in the schema: GRANT SELECT, UPDATE, INSERT, DELETE, INDEX, ALTER ON TABLE <schema.table> TO <user>

For more information about permissions required for DRS, see Here.

Network Settings

Plan the network properly to ensure that the source Oracle database can be accessed by UGO and DRS.

4.6.5 Creating a VPC and Security Group


Create a VPC and security group for a GaussDB instance.


Creating a VPC

- **Step 1** Log in to the management console.
- **Step 2** Click \bigcirc in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

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.

When its status becomes available, the VPC is created.


----End

Creating a Security Group


- **Step 1** Log in to the management console.
- **Step 2** Click \bigcirc in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and 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-01** in this example).
- **Step 9** Click the **Inbound Rules** tab, and then click **Add Rule**.



Step 10 Configure an inbound rule, add the IP address of the source database, and click **OK**.




----End


4.6.6 Creating a GaussDB Instance

This section describes how to create a GaussDB instance as the destination database.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Under the service list, choose **Databases** > **GaussDB**.
- **Step 4** In the navigation pane on the left, choose **GaussDB** > **Instances**.
- Step 5 Click Buy DB Instance.
- **Step 6** Configure the instance name and basic information.



Step 7 Configure instance specifications.




Select small specifications for the test. You are advised to configure specifications based on service requirements in actual use.

Step 8 Select a VPC and security group (created in **Creating a VPC and Security Group**) for the instance and configure the database port.



Step 9 Configure password and other information.




- Step 10 Click Next, confirm the information, and click Submit.
- **Step 11** Go to the instance list.

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

- **Step 12** Create a database in the destination GaussDB instance.
 - In the GaussDB instance list, locate the DB instance and click Log In in the Operation column.



- 2. In the displayed dialog box, enter the password and click **Test Connection**.
- 3. After the connection test is successful, click **Log In**.



Home DB Instance Name: gauss-c1b3 DB Engine Version: GaussDB Database List Database Name Table Quantity \$ Create Database * Name Database Name Only user databases can be created Character Set Template ① template0 Collation (?) Ctype ③ DBCOMPATIBILITY 3 Cancel

4. After the login is successful, click **Create Database**, enter required information as needed, and click **OK**.

5. Prepare permissions for logging in to GaussDB by referring to **Preparations**.


----End

4.6.7 Creating a UGO Task

4.6.7.1 Creating a Database Evaluation Project

- Step 1 Log in to the UGO console.
- **Step 2** In the navigation pane on the left, choose **Schema Migration** > **DB Evaluation**.
- **Step 3** Click **Create Project** in the upper right corner.
- **Step 4** Read **Source Database Preparation and Authorization Tips** and click **Create**.
- **Step 5** Configure the basic information.

After the basic information is configured, the **Test** button next to **Test Connection** is available.



Step 6 Click **Test** next to the **Test Connection** field.

If the connection test succeeds, the **Next** button will be available.


If the connection test fails, the message "Unable to connect to DB" is displayed.

- **Step 7** (Optional) Click **Test** next to **Network Stability**. A successful network stability test only means that there is little network latency or packet loss, or no packet loss at the current time. It takes 10s to 15s to complete.
- **Step 8** Click **Next** to go to the **Precheck** page.


The check result of each check item is displayed. You can also click **Recheck** to check the permissions again.



- **Step 9** After all check items are passed, click **Next** to go to the **Evaluation Scope Selection** page.
- **Step 10** Select the object types to be collected, target database versions, and source database schemas.



Step 11 Click **Next** to go to the **Confirmation** page.



- **Step 12** Confirm the settings and click **Create**. A message is displayed, indicating that the project is created successfully.
- **Step 13** Click **OK** to go to the **DB Evaluation** page. View the evaluation project you created in the list.




- Step 14 When Project Status is In progress. Confirm Target DB Pending, click Confirm Target DB Pending.
- Step 15 On the Target DB Analysis page, select GaussDB Primary/Standby-2.7 (select target database as needed) as the target database, and click Confirm DB Selection.
- **Step 16** In the displayed dialog box, click **Confirm**.
- Step 17 After the target database is confirmed, a dialog box is displayed. Click:Create Now to go to the page for migration project creation.

Create Later to stay on the current page.

----End

4.6.7.2 Creating an Object Migration Project

- Step 1 Log in to the UGO console.
- **Step 2** In the navigation pane on the left, choose **Schema Migration** > **Object Migration**.
- **Step 3** Click **Create Project** in the upper right corner.
- **Step 4** Configure required information.



Step 5 Click Test Connection.

If the connection test is successful, the **Create** button is available.

If the connection test fails, an error message is displayed.

Step 6 Click **Next** to go to the **Precheck** page.



Step 7 After the pre-check items are passed, click **Next** in the lower right corner to confirm the settings.



Step 8 Confirm the settings and click **Create**. On the object migration list page, view that the **Project Status** of the project is **Ready**.



----End

4.6.8 Creating a DRS Task

This section describes how to create a DRS task. Before creating a DRS task, you need to understand the prerequisites, suggestions, and precautions. For details, see **Precautions**.

Procedure

- **Step 1** Log in to the management console.
- Step 2 Click [♥] in the upper left corner and select a region.

 Select the region where the destination instance is deployed.
- **Step 3** Under the service list, choose **Databases** > **Data Replication Service**.
- **Step 4** In the navigation pane on the left, choose **Data Synchronization Management**. On the displayed page, click **Create Synchronization Task**.
- **Step 5** Configure synchronization instance information.
 - 1. Select a region, and project, and enter a task name.



2. Specify Data Flow, Source DB Engine, Destination DB Engine, Network Type, Destination DB Instance, Synchronization Instance Subnet (optional), Synchronization Mode, Specifications, AZ, and Tags (optional).

Figure 4-33 Synchronization instance details



3. Click **Create Now** and wait until the task status changes to **Configuring**.



----End


4.6.9 Starting Synchronization

4.6.9.1 UGO Schema Migration Phase 1: Tables, Primary Keys, Unique Keys/Indexes


Step 1 On the **Object Migration** page, locate the project that you created in **Creating an Object Migration Project** and click **Migrate** in the **Operation** column.




Step 2 On the **Conversion Plan** page, the collection objects and types for the project are displayed.




Step 3 Click **Next** to go to the **Conversion Config** page. Locate a feature and click **Edit** in the **Operation** column.




Step 4 Click Next to go to the Syntax Conversion page. Click Start to convert the syntax.



Step 5 Click Next to go to the Object Correction page. Select the objects to be skipped. In this phase, users, roles, permissions, tables, primary keys, unique keys, unique indexes, types, type bodies, functions, procedures, package bodies, and views can be migrated. You need to select objects that are not migrated and click Skip Migration.




For example, if the target database supports foreign keys, you can skip migrating them. In the object tree on the left, click **FOREIGN KEY**, select the FOREIGN KEY objects on the right, and click **Skip Migration**.



Step 6 Click **Next** to go to the **Migration & Verification** page. Click **Start** on the right to start the migration. Ensure that all objects except ignored objects are successfully migrated.

□ NOTE

Do not click **Finish** in the lower right corner because other objects need to be migrated later.




----End

4.6.9.2 DRS Synchronization Phase 1: Full Synchronization

- **Step 1** Select the full+incremental task created in section **Creating a DRS Task** and configure source and destination database information.
 - 1. Enter the IP address, port number, username, and password of the source database.


Click Test Connection.



Enter the username and password of the destination database. Click **Test Connection**.




3. Click **Next**. In the displayed box, read the message carefully and click **Agree**.




Step 2 Set the synchronization task.


1. Select the databases and tables of the source database to be synchronized. In this practice, all tables in **test** are selected.




2. Locate the database and table, respectively, and click **Edit** to change the database name and table name.



- 3. Click Next.
- **Step 3** On the **Advanced Settings** page, in the **Full Synchronization Settings** area, only select **Data** for **Synchronization Object Type**.




Step 4 (Optional) Process data based on the site requirements.




Step 5 After all settings are complete, perform a pre-check to ensure that the synchronization is successful. If any check fails, review the cause and rectify the fault. Then, click **Check Again**.

If all check items are successful, click Next.




Step 6 Confirm the task.

1. Check whether all configured information is correct.



- 2. Click Submit. In the display box, select I have read the precautions.
- 3. Click Submit.




Step 7 After the task is submitted, view and manage it.

1. After the task is created, return to the task list to view the task status. After the status changes to **Incremental**, click **Pause**.



2. In the displayed dialog box, deselect **Pause log capturing** and click **Yes**.



4.6.9.3 UGO Schema Migration Phase 2: Common Indexes

Step 1 Select the task in UGO Schema Migration Phase 1: Tables, Primary Keys, Unique Keys/Indexes. On the Object Correction page, select the indexes to be migrated (or click INDEX in the object tree on the left and select displayed indexes on the right), and click Undo Skip. Migration Status of the selected indexes changes to Manual.



Step 2 Click **Next** to go to the **Migration & Verification** page. Click **Start** to migrate indexes.

Ensure that all indexes are successfully migrated.

■ NOTE

Do not click **Finish** in the lower right corner because other objects need to be migrated later.



----End

4.6.9.4 DRS Synchronization Phase 2: Incremental Synchronization

Step 1 Locate the task paused in **DRS Synchronization Phase 1: Full Synchronization** and click **Resume** in the **Operation** column.




Step 2 The task enters the incremental synchronization phase. If the incremental synchronization delay is less than 30s, all data is about to be synchronized. Disconnect the service system of the source Oracle database to ensure that no data is written to the source Oracle database. Wait for about 1 minute after the disconnection.



- **Step 3** Compare data. Click the task name and choose **Synchronization Comparison**.
 - 1. On the **Object-Level Comparison** tab, click **Compare** to compare synchronization objects in the source Oracle database and destination GaussDB database.



2. On the **Data-Level Comparison** tab, click **Create Comparison Task**, select the tables whose rows to be compared, and click **OK**.




3. After the comparison task is complete, click **View Results** to view the row comparison details.



- **Step 4** After confirming that the data is consistent in the source and destination databases, locate the task and click **Stop** in the **Operation** column to stop the DRS task.
 - ----End

4.6.9.5 UGO Schema Migration Phase 3: Triggers, Events, Tasks, Foreign Keys, and Sequences

Step 1 Select the task in UGO Schema Migration Phase 2: Common Indexes. On the Object Correction page, select TRIGGER, EVENT, TASK, and FOREIGN KEY objects and click Undo Skip. Migration Status of the selected objects changes to Manual.



For FOREIGN KEY objects, select **FOREIGN KEY** in the object tree on the left and select displayed foreign keys on the right.



Step 2 Click **Next** to go to the **Migration & Verification** page. Click **Start** to migrate the selected objects.



Step 3 After the migration is successful, connect the service system to the target database GaussDB for data writing.

----End

4.7 From On-Premises Oracle to DDM

4.7.1 Overview

Description

You can use real-time synchronization of DRS to synchronize data from onpremises Oracle databases to Huawei Cloud DDM. Full+incremental synchronization can ensure that data is always in sync between the source Oracle database and the destination DDM instance.

Prerequisites

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

Service List


Virtual Private Cloud (VPC)

- Distributed Database Middleware (DDM)
- Data Replication Service (DRS)
- Data Admin Service (DAS)

Deployment Architecture

In this example, the source database is an on-premises Oracle database, and the destination database is a Distributed Database Middleware (DDM) instance on Huawei Cloud. DRS synchronizes data from the source database to the destination database through the public network. For details about the deployment architecture, see Figure 4-34.

Figure 4-34 Synchronizing data through the public 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.

4.7.2 Resource Planning

Table 4-25 Resource planning

Categor y	Subcategor y	Plan	Description
VPC	VPC name	vpc-DRStar	Specify a name that is easy to identify.
	Region	Test region	To achieve lower network latency, select the region nearest to you.
	AZ	AZ 3	-

Categor y	Subcategor y	Plan	Description
	Subnet	192.168.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-drs01	Specify a name that is easy to identify.
Oracle (source databas	Name	orcl	Specify a name that is easy to identify.
e)	Specificatio ns	16 vCPUs 32 GB	-
	Database version	11.2.0.1	-
	Database user	test_info	You can customize a user with the minimum permissions. For details, see here .
DDM (destinat ion	DDM instance name	ddm-DRS- AUTOTEST-001	Specify a name that is easy to identify.
databas e)	Region	Test region	To achieve lower network latency, select the region nearest to you.
	AZ	AZ 3	You can select one or more AZs. You are advised to create the instance across different AZs to improve service reliability.
	Node specification s	General-enhanced 4 vCPUs 8 GB	-
	Nodes	1	A single node has high availability risks. In practice, you are advised to create at least two nodes.
RDS instance associat ed with the DDM instance	RDS instance name	rds_ddm01	Specify a name that is easy to identify.
	Region	Test region	To achieve lower network latency, select the region nearest to you.
	Database version	MySQL 5.7	-

Categor y	Subcategor y	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/standby RDS instance and then deploy them in different AZs.
	Specificatio ns	General-purpose 4 vCPUs 8 GB	-
DRS synchron ization task	Synchroniza tion task name	DRS-OracleToDDM	Specify a name that is easy to identify.
	Source DB engine	Oracle	In this example, the source database is an Oracle database.
	Destination DB engine	DDM	In this example, the destination database is a DDM instance.
	Network type	Public network	Public network is used in this example.

4.7.3 Operation Process

Figure 4-35 shows the process of creating a DDM instance and synchronizing data from an Oracle database to DDM.

results




Figure 4-35 Flowchart

4.7.4 Preparing for the Source Oracle Database

Before the synchronization, prepare some data types in the source database for verification after the synchronization is complete.

The following table lists data types supported by DRS.

Confirm the task execution

results.

Table 4-26 Data type mapping

Data Type (Oracle)	Condition	Data Type (DDM)	Whether to Support Mapping
CHAR(n)	n ≤ 255	CHAR(n)	Yes
CHAR(n)	n > 255	VARCHAR(n)	Yes
VARCHAR(Size)	Length (row size) ≤ 65535	VARCHAR(n)	Yes
VARCHAR(Size)	Length (row size) > 65535	TEXT	Yes
VARCHAR2(n)	-	VARCHAR(n)	Yes
NCHAR(n)	n ≤ 255	NCHAR(n)	Yes
NCHAR(n)	n > 255	NVARCHAR(n)	Yes
NVARCHAR2(n)	-	NVARCHAR(n)	Yes
NUMBER(p,s)	s > 0	NUMBER(p,s)	Yes
NUMBER(p,s)	s ≤ 0	NUMBER(p-s,0)	Yes
BINARY_FLOAT	-	FLOAT	Yes
BINARY_DOUBL E	-	DOUBLE	Yes
FLOAT(b)	b ≤ 99	DECIMAL(b*0.30103*2, b*0.30103)	Yes
FLOAT(b)	b > 99	DOUBLE	Yes
DATE	-	DATETIME	Yes
TIMESTAMP	-	TIMESTAMP	Yes
TIMESTAMP WITH LOCAL TIME ZONE	-	TIMESTAMP	Yes
TIMESTAMP WITH TIME ZONE	-	TIMESTAMP	Yes
INTERVAL	Incremental	VARCHAR(30)	No
INTERVAL	Full; 6 digit precision	VARCHAR(30)	Yes
BLOB	-	LONGBLOB	Yes
CLOB	-	LONGTEXT	Yes
NCLOB	-	LONGTEXT	Yes

Data Type (Oracle)	Condition	Data Type (DDM)	Whether to Support Mapping
LONG	-	LONGTEXT	Yes
LONG_RAW	-	LONGBLOB	Yes
RAW	-	VARBINARY	Yes
ROWID	-	VARCHAR(18)	Yes

Perform the following steps to generate data in the source database:

- **Step 1** Use a database connection tool to connect to the source Oracle database based on its IP address.
- **Step 2** Construct data in the source database based on its supported data types.
 - 1. Create a test user.

create user test_info identified by xxx;

test_info indicates the user created for the test, and *xxx* indicates the password of the user.

- 2. Assign permissions to the user. **grant dba to** *test info*;
- 3. Create a data table under the user.

```
CREATE TABLE test_info.table3(
ID INT,
COL01 CHAR(100),
COL02 NCHAR(100),
PRIMARY KEY(ID)
).
```

4. Insert two rows of data.

```
insert into test_info.table3 values(4,'huawei','xian');
insert into test_info.table3 values(2,'DRS-test','test1');
insert into test_info.table3 values(1,'huawei','xian');
```

5. Make the above statements take effect.

commit;

----End

4.7.5 Preparing for the Destination DDM Instance

4.7.5.1 Creating a VPC and Security Group

Create a VPC and security group to prepare for the destination DDM instance.


Creating a VPC

Step 1 Log in to the management console.

- **Step 2** Click oin the upper left corner and select a region and project.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

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.

When its status becomes available, the VPC is created.


----End

Creating a Security Group

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.
- **Step 3** Click the service list icon on the left and 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** Specify the security group name and other information, and click **OK**.



----End

4.7.5.2 Creating a DDM Instance

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the displayed page, in the upper right corner, click **Buy DDM Instance**.
- **Step 5** Specify the instance information and specifications as required.



Step 6 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 7** After the configuration is complete, click **Next** at the bottom of the page.
- **Step 8** Confirm your settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Submit.
- **Step 9** Go to the **Instances** page to view and manage the instance.

The default database port is **5066** and can be changed after a DDM instance is created. If the status of the instance is **Running**, the instance has been created.

----End

4.7.5.3 Creating an RDS for MySQL Instance


Create an RDS for MySQL instance associated with the DDM instance.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and 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.



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**.



The RDS for MySQL instance must be in the same VPC and subnet as your DDM instance.



Step 8 Configure the instance password.



Step 9 Click Next.

- **Step 10** Confirm your settings.
 - To modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**.
- **Step 11** Return to the instance list.

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

----End


4.7.5.4 Creating a Schema and Associating It with the RDS for MySQL Instance

- Step 1 Log in to the management console.
- **Step 2** Click in the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the **Instances** page, locate the required DDM instance and click **Create Schema** in the **Operation** column.
- **Step 5** On the displayed page, specify a sharding method, enter a schema name, set the number of shards, select the required DDM accounts, and click **Next**.

In this example, the schema is unsharded, and the schema name is **db_test**.

A CAUTION


Currently, only data can be synchronized from Oracle to DDM. To synchronize table structures and other objects, you need to create schemas in the destination DDM instance based on table structures of the source Oracle database.



----End

4.7.5.5 Creating a DDM Account

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the **Instances** page, locate the required DDM instance and click its name.
- **Step 5** In the navigation pane, choose **Accounts**.
- **Step 6** On the displayed page, click **Create Account** and configure parameters as needed. For details about the permissions required by the DDM account, see **Precautions**



Step 7 Click OK.

----End

4.7.5.6 Creating Table Structures in the Destination Database

Currently, only data can be synchronized from the source Oracle database to the destination DDM instance. Source table structures and other objects cannot be synchronized, so you need to create table structures and indexes in the destination database based on the table structures of the source schema. For more details, see **Precautions**.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and choose Databases > Distributed Database Middleware.
- **Step 4** On the displayed page, enter the username and password created in **Creating a DDM Account**, and click **Test Connection**.
- **Step 5** After the connection is successful, click **Log In** to log in to the DDM instance.
- Step 6 Click the db_test schema created in Creating a Schema and Associating It with the RDS for MySQL Instance.
- **Step 7** Run the following SQL statements in database **db_test** to create table **table3** with the same structure as the source:

CREATE TABLE `db_test`.`table3`(
ID INT,
COL01 CHAR(100),
COL02 NCHAR(100),
PRIMARY KEY(ID));

----End

4.7.6 Creating a DRS Synchronization Task

This section describes how to create a DRS task and synchronize data from an onpremises Oracle database to Huawei Cloud DDM.

Pre-Check

Before creating a task, check whether synchronization conditions are met.

In this example, data is synchronized to DDM. For more details, see **Precautions**.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.

- **Step 3** Click the service list icon on the left and choose **Databases** > **Data Replication Service**.
- **Step 4** In the upper right corner, click **Create Synchronization Task**.
- **Step 5** Configure synchronization task parameters.
 - 1. Specify a synchronization task name.



2. Enter the synchronization task information and select the destination database.

Select the DDM instance created in **Creating a DDM Instance** as the destination database.


Figure 4-36 Synchronization instance details



Step 6 Click Create Now.

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

- **Step 7** Configure source and destination database information.
 - 1. Configure the source database information and click **Test Connection**. If a successful test message is returned, the database is connected.



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




Step 8 Click Next.

Step 9 On the **Set Synchronization Task** page, select the synchronization objects.

- Set Flow Control to No.
- Set Synchronization Object to Tables.

Select the databases and tables of the source database to be synchronized. In this best practice, **table3** in **test_info** is selected and synchronized to **table3** in **db_test**.




Step 10 Click **Next**. 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**.
- If all check items are successful, click Next.

Step 11 Click Submit.

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

It takes several minutes to complete.



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

- For synchronization from Oracle to DDM, full synchronization and full+incremental synchronization modes are supported.
- If you create a full synchronization task, the task automatically stops after the full data is synchronized to the destination.
- If you create a full+incremental synchronization task, a full synchronization is executed first. After the full synchronization is complete, an incremental synchronization starts.
- During the incremental synchronization, data is continuously synchronized, so the task will not automatically stop.

----End

4.7.7 Confirming Synchronization Results

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

- DRS compares synchronization objects and data and provides comparison results. For details, see Viewing Synchronization Results on the DRS Console.
- 2. Log in to the destination side to check whether the databases, tables, and data are synchronized. Confirm the data synchronization status. For details, see Viewing Synchronization Results on the DDM Console.


Viewing Synchronization Results on the DRS Console

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and choose Databases > Data Replication Service.
- **Step 4** Locate the required DRS instance and click its name.
- **Step 5** In the navigation pane on the left, choose **Synchronization Comparison**.
- **Step 6** Click the **Object-Level Comparison** tab and check whether some objects are missing.

Click **Compare**. After the comparison is complete, view the comparison results.



- **Step 7** Click the **Data-Level Comparison** tab and check whether the number of rows of synchronized objects is consistent.
 - 1. Click Create Comparison Task.
 - 2. In the displayed dialog box, select the comparison type, time, and object.



3. After the comparison task is complete, view the data comparison results.



4. To view the comparison details, click **View Results** next to the comparison task.



----End

Viewing Synchronization Results on the DDM Console

- **Step 1** Log in to the management console.
- **Step 2** Click $^{\bigcirc}$ in the upper left corner and select a region and project.
- Step 3 Click the service list icon on the left and choose Databases > Distributed Database Middleware.
- **Step 4** Locate the target DDM instance and click **Log In** in the **Operation** column.
- **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** View destination database and table names and check whether all data is synchronized.

----End

4.8 From RDS for MySQL to Kafka

4.8.1 Overview

Description

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

- How to create an RDS for MySQL instance.
- How to create DMS for Kafka.
- 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 for 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 4-37**.

Figure 4-37 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.

4.8.2 Resource and Cost Planning

Table 4-27 Resource planning

Categ ory	Subcatego ry	Plan	Description
VPC	VPC name	vpc-DRStest	Specify a name that is easy to identify.
	Region	AP-Singapore	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 (sourc e datab ase)	RDS instance name	rds-mysql	Specify a name that is easy to identify.
	DB engine version	MySQL 5.7	-

Categ ory	Subcatego ry	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.
	Specificati ons	General-purpose 4 vCPUs 8 GB	-
Kafka (desti nation	Kafka instance name	kafka-drs	Specify a name that is easy to identify.
datab ase)	Version	2.3.0	-
asey	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.
	Specificati ons	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 synchr onizati on task	Synchroniz ation task name	DRS- MySQLToKafka	Custom
	Source DB engine	MySQL	In this example, the source is an RDS for MySQL instance.
	Destinatio n DB engine	Kafka	In this example, the destination database is Kafka.
	Network Type	VPC	In the practice, select the VPC network.

□ 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.

4.8.3 Operation Process

Figure 4-38 shows the process of creating an RDS for MySQL instance and synchronizing the incremental data from an RDS for MySQL instance to Kafka.




Figure 4-38 Flowchart


4.8.4 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 AP-Singapore.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**. The VPC console is displayed.

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 AP-Singapore.
- 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**.



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




----End

4.8.5 Preparing for Source RDS for MySQL


4.8.5.1 Creating an RDS for MySQL Instance

Create an RDS for 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 AP-Singapore.
- **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.



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 information such as the instance password and parameter template.



- **Step 9** Click **Next**, confirm the information, and click **Submit**.
- **Step 10** Return to the instance list.


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

----End

4.8.5.2 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 AP-Singapore.
- **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** Click **Create Database** to create the **db_test** database.



Step 8 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;
  SQL Preview
     1 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;
                                              Back
```

----End

4.8.6 Preparing for Destination Kafka


4.8.6.1 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 AP-Singapore.
- Step 3 Click the service list icon on the left and choose Middleware > Distributed Message Service (for Kafka).
- Step 4 Click Buy 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 Buy, confirm the information, and click Submit.
- Step 11 Return to the instance list.

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

----End

4.8.6.2 Creating a Topic

- Step 1 Click a Kafka instance.
- **Step 2** Choose **Topics**, and click **Create Topic**.
- **Step 3** In the dialog box that is displayed, enter a topic name, specify other parameters, and click **OK**.



----End

4.8.7 Creating a DRS Synchronization Task

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

Pre-Check

Before creating a task, check the synchronization conditions.

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


Procedure

- Step 1 Log in to the management console.
- **Step 2** Click in the upper left corner of the management console and select AP-Singapore.
- **Step 3** Under the service list, choose **Databases** > **Data Replication Service**.
- **Step 4** Choose **Data Synchronization Management** and click **Create Synchronization Task**.
- **Step 5** Configure synchronization task parameters.
 - 1. Task name




Select the source database, destination database, and network information.
 Select the RDS instance created in Preparing for Source RDS for MySQL as the source database.

Figure 4-39 Synchronization instance details



3. Set Enterprise Project to default.




Step 6 Click Create Now.

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, the database is connected.

Source Database



- 3. Select the VPC and subnet where the destination database is located, and enter the Kafka IP address and port number.
- 4. Click **Test Connection**.

If a successful test message is returned, the database is connected.




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 4-28 Sync Settings

Туре	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.

Step 11 Select a data processing mode. Data synchronization from RDS for 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**, the synchronization task has been started.

■ NOTE

- In this example, **Synchronization Mode** is set to **Incremental** for the task from RDS for MySQL to Kafka. After the task is started, the status is **Incremental**.
- 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

4.8.8 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 for 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 AP-Singapore.
- **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** 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');



- **Step 8** In the Service List, choose **Application** > **Distributed Message Service for Kafka** to open the Kafka console.
- **Step 9** Click a Kafka instance.
- **Step 10** Click the **Message Query** tab, select the Kafka topic, view the received data in JSON format.




Step 11 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**.



----End

4.9 From GaussDB Distributed to Kafka

Description

This practice describes how to create a DRS synchronization task to synchronize incremental data from GaussDB Distributed to Kafka.

Prerequisites

- You have registered with Huawei Cloud.
- Your account balance is greater than or equal to \$0 USD.
- You have logged in to the DRS console.


Service List

- Virtual Private Cloud (VPC)
- GaussDB Distributed

- DMS for Kafka
- Data Replication Service (DRS)

Architecture of a Distributed GaussDB Instance

Figure 4-40 Logical architecture of a distributed GaussDB instance



GaussDB consists of Global Transaction Manager (GTM), Operation Manager (OM), Cluster Manager (CM), Coordinator Node (CN), and Data Node (DN). An application sends a SQL query request to a CN. The SQL statement can be INSERT, DELETE/DROP, UPDATE, or SELECT. The CN uses the optimizer of the database to generate an execution plan and sends the plan to DNs. Each DN processes data based on the execution plan. After the processing is complete, DNs return the result set to the CN for summary. Finally, the CN returns the summary result to the application.

You can add nodes for distributed instances as needed to handle large volumes of concurrent requests. For details, see **Adding Nodes**.

DRS synchronization network diagram

In this example, the source is a distributed GaussDB instance, and the destination is a DMS for Kafka instance. Incremental data of the source database is synchronized to the destination database through a VPC. The following figure shows the deployment architecture.

DN 1 Connecting the Connecting the DRS destination database source database Subtask 1 Huawei Cloud through a VPC GaussDB -Distributed Connecting the DMS for Kafka destination database through a VPC DN₂ Connecting the DRS source database

Subtask 2

Figure 4-41 Data synchronization from GaussDB Distributed to Kafka through a

Notes on Usage

- 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

through a VPC

Resource Planning


Table 4-29 Resource planning

Category	Subcate gory	Planned Value	Description
VPC	VPC name	vpc-DRStest	Specify a name that is easy to identify.
	Subnet CIDR block	10.0.0.0/24	Select a subnet with sufficient network resources.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	Subnet name	subnet-drs01	Specify a name that is easy to identify.
GaussDB (source	Instance name	drs-gaussdbv5- src-1	Specify a name that is easy to identify.
database)	DB engine version	GaussDB 8.102	-
	DB instance type	Distributed	Select a proper instance type based on its description.
	Storage type	Ultra-high I/O	GaussDB supports ultra-high I/O storage with a maximum throughput of 350 MB/s.

Category	Subcate gory	Planned Value	Description
	Specifica tions	Dedicated 8 vCPUs 64 GB	Select the specifications as needed.
Kafka (destinatio n	Kafka instance name	kafka-drs	Specify a name that is easy to identify.
database)	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.
	Specifica tions	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 synchroniz ation task	Synchron ization task name	DRS- GaussDBToKafka	Specify a name that is easy to identify.
	Source DB engine	GaussDB Distributed	-
	Destinati on DB engine	Kafka	In this example, the destination database is a Kafka instance.
	Network type	VPC	When creating a task, select VPN or Direct Connect.

Flowchart

The following figure shows the process of creating a DRS task and synchronizing the incremental data from a distributed GaussDB instance to Kafka.




Creating a VPC

Create a VPC to prepare network resources for creating a GaussDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

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

Create a security group for creating a GaussDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click oin the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and 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.

----End


Creating a Distributed GaussDB Instance

The following describes how to create a distributed GaussDB instance as the destination database.

- Step 1 Log in to the management console.
- **Step 2** Click on the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Databases > GaussDB.
- Step 4 Click Buy DB Instance.
- **Step 5** Configure the instance name and basic information.



Step 6 Configure instance specifications.




Select small specifications for this test instance. You are advised to configure specifications based on service requirements in actual use.

Step 7 Select a VPC and security group (created in **Creating a VPC** and **Creating a Security Group**) for the instance and configure the database port.



Step 8 Configure password and other information.




- **Step 9** Click **Next**, confirm the information, and click **Submit**.
- Step 10 Go to the instance list.

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

----End

Constructing Test Data for the GaussDB Instance

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Databases** > **GaussDB**.
- **Step 4** Select the target GaussDB 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 Click Create Database to create the db_test database.



Step 8 Run the following statement in **db_test** to create table **schema_test.table1**:

create table schema_test.table1(c1 int primary key,c2 varchar(10),c3 TIMESTAMP(6));

----End


Creating a Kafka Instance

The following describes how to create a Kafka instance.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Middleware > Distributed Message Service (for Kafka).
- Step 4 Click Buy 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 a VPC and security group (created in **Creating a VPC** and **Creating a Security Group**) for the instance.



Step 9 Configure the instance password.




- Step 10 Click Buy, confirm the information, and click Submit.
- **Step 11** Go to the instance list.

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

----End

Creating a Kafka Topic

- **Step 1** Click a Kafka instance.
- **Step 2** Choose **Topics**, and click **Create Topic**.
- **Step 3** In the dialog box that is displayed, enter a topic name, specify other parameters, and click **OK**.



----End

Creating a DRS Synchronization Task

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Databases > Data Replication Service.
- **Step 4** Choose **Data Synchronization Management** and click **Create Synchronization Task**.
- **Step 5** Configure synchronization task parameters.
 - 1. Specify a synchronization task name.



Select the source database, destination database, and network information.
 Select the GaussDB instance created in Creating a Distributed GaussDB Instance as the source database.

Synchronization Instance Details ©

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

* Data Prov

**Data Pr

Figure 4-42 Synchronization instance details

3. Select specifications and AZ.



Step 6 Click Create Now.

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, the database is connected.

Source Database



- 3. Select the VPC and subnet where the destination database is located, and enter the Kafka IP address and port number.
- 4. Click Test Connection.

If a successful test message is returned, the database is connected.



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 4-30 Synchronization settings

Туре	Setting
Topic Synchronization Policy	Deliver the content to a topic named testTopic . Different topic policies correspond to different Kafka partition policies. For details, see Topic and Partition Synchronization Policies .
Synchronize Topic To	Partition 0 Different topic policies correspond to different Kafka partition policies. For details, see Topic and Partition Synchronization Policies.
Data Format in Kafka	You can select the JSON format. For details, see Kafka Message Format .
Synchronization Object	Select schema_test.table1 in the db_test database.

- **Step 10** Click **Next** and wait for the check results.
- **Step 11** If the check is complete and the check success rate is 100%, click **Next**.
- **Step 12** 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**, the synchronization task has been started.

- In this example, Synchronization Mode is set to Incremental for the task from GaussDB Distributed to Kafka. After the task is started, the status is Incremental.
- If you create a full+incremental synchronization task, a full synchronization is executed first. After the full synchronization is complete, an incremental synchronization starts.
- During the incremental synchronization, data is continuously synchronized, so the task will not automatically stop.

----End

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 GaussDB database and viewing the data received by Kafka.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Databases** > **GaussDB**.
- **Step 4** Locate the required GaussDB 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** Run the following statement to insert data to the **db_test.schema_test.table1** table.

```
insert into schema_test.table1 values(1,'testKafka',current_timestamp(6)); update schema_test.table1 set c2 ='G2K' where c1 =1; delete schema_test.table1 where c1 =1;
```

Step 8 On the Kafka client, check the received data in JSON format.

./kafka-console-consumer.sh --bootstrap-server *ip:port* --topic testTopic --from-beginning



Step 9 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**.

----End

Topic and Partition Synchronization Policies

Table 4-31 Topic and partition policies

Topic Policy	Available Partition Policies	Description
A specified topic Select A specified topic if the data volume of the source database is small.	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
	Partitions are identified by the hash values of the primary key	This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
	Partitions are differentiated by the hash values of database_name.schema_name	This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently.

Topic Policy	Available Partition Policies	Description
	Partitions are differentiated by the hash values of database_name.DN_sequence_number	This mode applies to scenarios where multiple databases correspond to one topic, preventing data of multiple data nodes from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
	Partitions are differentiated by the hash values of <i>non-primary-key columns</i>	This mode is for when one table corresponds to one topic, preventing table data from being written to the same partition. Users can customize message keys based on the hash values of non-primary-key columns, and consumers can obtain data from different partitions concurrently.
Automatically generated using the database_name-schema_name-table_name format	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
If each table contains a lot of data, select Automatically generated using the database_name-schema_name-table_name format.	Partitions are identified by the hash values of the primary key	This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently.

Topic Policy	Available Partition Policies	Description
	Partitions are differentiated by the hash values of <i>non-primary-key columns</i>	This mode is for when one table corresponds to one topic, preventing table data from being written to the same partition. Users can customize message keys based on the hash values of non-primary-key columns, and consumers can obtain data from different partitions concurrently.
Automatically generated based on the database name If the source database does not contain a lot of data, select Automatically	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
generated based on the database name.	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
	Partitions are differentiated by the hash values of database_name.schema_name	This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
	Partitions are differentiated by the hash values of database_name.DN_sequence_number	This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple DNs from being written to the same partition, so that consumers can obtain data from different partitions concurrently.

Topic Policy	Available Partition Policies	Description
Automatically generated using the database_name- schema_name format If each schema contains a lot of data, select Automatically	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
generated using the database_name-schema_name format.	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
Automatically generated using the database_name-DN_sequence_number format If there is a large amount of data on each	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
data node, select Automatically generated using the database_name- DN_sequence_number format.	Partitions are differentiated by the hash values of database_name.schema_name	This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.

4.10 From GaussDB Centralized to Kafka

Description

This practice describes how to create a DRS synchronization task to synchronize incremental data from GaussDB Centralized to Kafka.

Prerequisites


- You have registered with Huawei Cloud.
- Your account balance is greater than or equal to \$0 USD.
- You have logged in to the DRS console.

Service List

- Virtual Private Cloud (VPC)
- GaussDB Centralized
- DMS for Kafka
- Data Replication Service (DRS)

Architecture of a Primary/Standby GaussDB Instance

Figure 4-43 Logical architecture of a primary/standby GaussDB instance



GaussDB Centralized consists of Operation Manager (OM), Cluster Manager (CM), and Data Node (DN). An application sends a task directly to the DN, and the DN returns the result to the application after processing the task.

Primary/Standby instances are suitable for scenarios with small and stable volumes of data, where data reliability and service availability are extremely important.

DRS synchronization network diagram

In this example, the source is a centralized GaussDB instance, and the destination is a DMS for Kafka instance. Incremental data of the source database is synchronized to the destination database through a VPC. The following figure shows the deployment architecture.

Figure 4-44 Data synchronization from GaussDB Centralized to Kafka through a VPC



Notes on Usage

- 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.

Resource Planning


Table 4-32 Resource planning

Category	Subcate gory	Planned Value	Description
VPC	VPC name	vpc-DRStest	Specify a name that is easy to identify.
	Subnet CIDR block	10.0.0.0/24	Select a subnet with sufficient network resources.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	Subnet name	subnet-drs01	Specify a name that is easy to identify.
GaussDB (source database)	Instance name	drs-gaussdb-src-1	Specify a name that is easy to identify.
	DB engine version	GaussDB 8.1.0	-
	DB instance type	Primary/standby	Select a proper instance type based on its description.
	Storage type	Ultra-high I/O	GaussDB supports ultra-high I/O storage with a maximum throughput of 350 MB/s.
	Specifica tions	General-purpose 4 vCPUs 16 GB	Select the specifications as needed.

Category	Subcate gory	Planned Value	Description
Kafka (destinatio n	Kafka instance name	kafka-drs	Specify a name that is easy to identify.
database)	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.
	Specifica tions	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 synchroniz ation task	Synchron ization task name	DRS- GaussDBToKafka	Specify a name that is easy to identify.
	Source DB engine	GaussDB Centralized	-
	Destinati on DB engine	Kafka	In this example, the destination database is a Kafka instance.
	Network type	VPC	When creating a task, select VPN or Direct Connect .

Flowchart

The following figure shows the process of creating a DRS task and synchronizing the incremental data from a centralized GaussDB instance to Kafka.




Creating a VPC

Create a VPC to prepare network resources for creating a GaussDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click oin the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Networking** > **Virtual Private Cloud**.

The VPC console is displayed.

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

Create a security group for creating a GaussDB instance.

- **Step 1** Log in to the management console.
- **Step 2** Click oin the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and 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.

----End


Creating a Primary/Standby GaussDB Instance

The following describes how to create a primary/standby GaussDB instance as the destination database.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Databases > GaussDB.
- Step 4 Click Buy DB Instance.
- **Step 5** Configure the instance name and basic information.



Step 6 Configure instance specifications.




Select small specifications for this test instance. You are advised to configure specifications based on service requirements in actual use.

Step 7 Select a VPC and security group (created in **Creating a VPC** and **Creating a Security Group**) for the instance and configure the database port.



Step 8 Configure password and other information.




- **Step 9** Click **Next**, confirm the information, and click **Submit**.
- Step 10 Go to the instance list.

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

----End

Constructing Test Data for the GaussDB Instance

- Step 1 Log in to the management console.
- **Step 2** Click \bigcirc in the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Databases > GaussDB.
- **Step 4** Select the target GaussDB 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** Click **Create Database** to create the **db_test** database.



Step 8 Run the following statement in **db_test** to create table **schema_test.table1**:

create table schema_test.table1(c1 int primary key,c2 varchar(10),c3 TIMESTAMP(6));

----End


Creating a Kafka Instance

The following describes how to create a Kafka instance.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Middleware > Distributed Message Service (for Kafka).
- Step 4 Click Buy 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 a VPC and security group (created in **Creating a VPC** and **Creating a Security Group**) for the instance.



Step 9 Configure the instance password.




- Step 10 Click Buy, confirm the information, and click Submit.
- **Step 11** Go to the instance list.

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

----End

Creating a Kafka Topic

- **Step 1** Click a Kafka instance.
- **Step 2** Choose **Topics**, and click **Create Topic**.
- **Step 3** In the dialog box that is displayed, enter a topic name, specify other parameters, and click **OK**.



----End

Creating a DRS Synchronization Task

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select a region.
- Step 3 Click the service list icon on the left and choose Databases > Data Replication Service.
- **Step 4** Choose **Data Synchronization Management** and click **Create Synchronization Task**.
- **Step 5** Configure synchronization task parameters.
 - 1. Specify a synchronization task name.



 Select the source database, destination database, and network information.
 Select the GaussDB instance created in Creating a Primary/Standby GaussDB Instance as the source database.

Synchronization instance Details ©

The following information cannot be modified after you got to the red page.

© Data Place

To the cloud.

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

© Source DB Engine

MySQL

DDM

Gauss DB Detebuted

Council Developed

Developed Developed

Developed Developed

Developed Developed

Developed Developed Developed

Developed Developed Developed

Developed Developed Developed

Developed Developed Developed

Developed Developed Developed Developed

Developed Devel

Figure 4-45 Synchronization instance details

3. Select specifications and AZ.



Step 6 Click Create Now.

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, the database is connected.

Source Database



- 3. Select the VPC and subnet where the destination database is located, and enter the Kafka IP address and port number.
- 4. Click Test Connection.

If a successful test message is returned, the database is connected.



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 4-33 Synchronization settings

Туре	Setting
Topic Synchronization Policy	Deliver the content to a topic named testTopic .
	Different topic policies correspond to different Kafka partition policies. For details, see Topic and Partition Synchronization Policies.
Synchronize Topic To	Partition 0
	Different topic policies correspond to different Kafka partition policies. For details, see Topic and Partition Synchronization Policies.
Data Format in Kafka	You can select the JSON format. For details, see Kafka Message Format .
Synchronization Object	Select schema_test.table1 in the db_test database.

- **Step 10** Click **Next** and wait for the check results.
- **Step 11** If the check is complete and the check success rate is 100%, click **Next**.
- **Step 12** 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**, the synchronization task has been started.

- In this example, **Synchronization Mode** is set to **Incremental** for the task from GaussDB Primary/Standby to Kafka. After the task is started, the status is **Incremental**.
- During the incremental synchronization, data is continuously synchronized, so the task will not automatically stop.

----End

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 GaussDB database and viewing the data received by Kafka.

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select a region.
- **Step 3** Click the service list icon on the left and choose **Databases** > **GaussDB**.
- **Step 4** Locate the required GaussDB 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** Run the following statement to insert data to the **db_test.schema_test.table1** table.

```
insert into schema_test.table1 values(1,'testKafka',current_timestamp(6)); update schema_test.table1 set c2 ='G2K' where c1 =1; delete schema_test.table1 where c1 =1;
```

Step 8 On the Kafka client, check the received data in JSON format.

./kafka-console-consumer.sh --bootstrap-server *ip:port* --topic testTopic --from-beginning



Step 9 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.

----End

Topic and Partition Synchronization Policies

Table 4-34 Topic and partition policies

Topic Policy	Available Partition Policies	Description
A specified topic Select A specified topic if the data volume of the source database is small.	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
	Partitions are identified by the hash values of the primary key	This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
	Partitions are differentiated by the hash values of database_name.schema_name	This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently.

Topic Policy	Available Partition Policies	Description
	Partitions are differentiated by the hash values of <i>non-primary-key columns</i>	This mode is for when one table corresponds to one topic, preventing table data from being written to the same partition. Users can customize message keys based on the hash values of non-primary-key columns, and consumers can obtain data from different partitions concurrently.
Automatically generated using the database_name-schema_name-table_name format	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
If each table contains a lot of data, select Automatically generated using the database_name-schema_name-table_name format.	Partitions are identified by the hash values of the primary key	This mode applies to scenarios where one table corresponds to one topic, preventing table data from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
	Partitions are differentiated by the hash values of <i>non-primary-key columns</i>	This mode is for when one table corresponds to one topic, preventing table data from being written to the same partition. Users can customize message keys based on the hash values of non-primary-key columns, and consumers can obtain data from different partitions concurrently.

Topic Policy	Available Partition Policies	Description
Automatically generated based on the database name If the source database does not contain a lot of data, select Automatically	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
generated based on the database name.	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.
	Partitions are differentiated by the hash values of database_name.schema_name	This mode applies to scenarios where one database corresponds to one topic, preventing data of multiple schemas from being written to the same partition, so that consumers can obtain data from different partitions concurrently.
Automatically generated using the database_name- schema_name format If each schema contains a lot of data, select Automatically	Partitions are differentiated by the hash values of database_name.schema_name.table_name	This mode is recommended in single-table query scenarios where the read and write performance on the single table can be improved.
generated using the database_name-schema_name format.	Partition 0	If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted.

5 Real-Time Disaster Recovery

5.1 Configuring Remote Single-Active DR for an RDS for MySQL Instance Using DRS

5.1.1 Overview

Scenarios

This best practice involves two tasks:

- Create an RDS for MySQL instance.
- Use DRS to establish a remote single-active DR relationship for the RDS for MySQL instance.

Prerequisites

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

How Cross-Region DR Works

RDS for MySQL instances are deployed in the production and DR data centers. DRS replicates data from the production center to the DR center, keeping data synchronous between your primary instance and the DR instance.

Service List

- Virtual Private Cloud (VPC)
- Elastic IP (EIP)
- Relational Database Service (RDS)
- Data Replication Service (DRS)

Notes on Usage

- 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 RDS for MySQL instance DR, see From MySQL to MySQL (Single-Active DR).

5.1.2 Resource Planning

Table 5-1 Resource planning

Categor y	Subcategor y	Planned Value	Description
VPC in the	VPC name	vpc-01	Specify a name that is easy to identify.
producti on center	Region	CN-Hong Kong	To achieve lower network latency, select the region nearest to you.
	AZ	AZ2	-
	Subnet	192.168.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-3c29	Specify a name that is easy to identify.
VPC in the DR	VPC name	vpc-DR	Specify a name that is easy to identify.
center	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	AZ	AZ1	-
	Subnet	192.168.0.0/24	Select a subnet with sufficient network resources.
	Subnet name	subnet-ac27	Specify a name that is easy to identify.
RDS for MySQL	Instance name	rds-database-01	Specify a name that is easy to identify.
instance in the producti on	Region	CN-Hong Kong	To achieve lower network latency, select the region nearest to you.
center	DB engine version	MySQL 8.0	-

Categor y	Subcategor y	Planned Value	Description
	Instance type	Single	A single instance is used in this example. To improve service reliability,
			select a primary/standby instance.
	Storage type	Ultra-high I/O	-
	AZ	AZ2	AZ2 is selected in this example.
			To improve service reliability, select the primary/standby instance type and deploy the primary and standby instances in different AZs.
	Instance specification s	General-enhanced 2 vCPUs 4 GB	-
RDS for MySQL instance in the DR center	Instance name	rds-DR	Specify a name that is easy to identify.
	Region	AP-Singapore	To achieve lower network latency, select the region nearest to you.
	DB engine version	MySQL 8.0	-
	Instance type	Single	A single instance is used in this example.
			To improve service reliability, select a primary/standby instance.
	Storage type	Cloud SSD	-
	AZ	AZ1	AZ1 is selected in this example.
			To improve service reliability, select the primary/standby instance type and deploy the primary and standby instances in different AZs.
	Instance specification s	General-purpose 2 vCPUs 8 GB	-
DRS DR task	DR task name	DRS-DR-Task	Specify a name that is easy to identify.

Categor y	Subcategor y	Planned Value	Description
	Source DB engine	MySQL	In this example, the primary instance created in CN-Hong Kong is used as the source database.
	Destination DB engine	MySQL	In this example, the DR instance created in AP-Singapore is used as the destination database.
	Network type	Public network	Public network is used in this example.

5.1.3 Operation Process

You can create a single RDS instance and a DR instance and migrate data from the single instance to the DR instance.




Figure 5-1 Flowchart

5.1.4 Configuring an RDS for MySQL Instance in the Production Center


5.1.4.1 Creating a VPC and Security Group

Create a VPC and security group for a DB instance in the production center.

Creating a VPC

- **Step 1** Go to the **Create VPC** page.
- **Step 2** On that page, select **CN-Hong Kong** for **Region**, and configure the basic information, subnet, and IP address.

Figure 5-2 Creating a VPC




Step 3 Click Create Now.

----End

Creating a Security Group

- **Step 1** Log in to the management console.
- **Step 2** Click on the upper left corner and select **CN-Hong Kong**.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**.
- **Step 4** In the navigation pane on the left, choose **Access Control** > **Security Groups**.
- Step 5 Click Create Security Group.

Figure 5-3 Creating a security group



Step 6 Click Create Now.

----End


5.1.4.2 Creating an EIP

Create an EIP for your source DB instance. Using the EIP, external systems can access your application and DRS can connect to the source DB instance.

Procedure

- Step 1 Go to the Buy EIP page.
- **Step 2** On that page, select **CN-Hong Kong** for **Region**, and configure the basic information and bandwidth as prompted.

Figure 5-4 Buying an EIP



- Step 3 Click Next.
- **Step 4** Confirm the information and click **Submit**.
 - ----End

5.1.4.3 Creating an RDS for MySQL Instance

Create an RDS for MySQL instance (source database), and select the VPC and EIP you configured for the instance.

Procedure

- Step 1 Go to the Buy DB Instance page.
- Step 2 Select CN-Hong Kong for Region. Configure instance information and click Buy.





Figure 5-5 Selecting a DB engine


Figure 5-6 Selecting specifications


cn-north-4a AZ7 cn-north-4b cn-north-4c



Basic Settings DB Instance Name (?) If you buy multiple DB instances at a time, they will be named with four digits appended in the format "DB instance name-SN". For example, if the DB instance name is "instance", the first instance will be named "instance-0001", the second "instance-0002", Password Skip Configure To log in, you will have to reset the password later on the Basic Information page for this instance. VPC (?) ✓ Q Create VPC 🗹 The VPC an RDS instance is deployed in cannot be changed later. ECSs in different VPCs cannot communicate with each other by default. subnet-a389(192.168.0.0/24) IPv6 CIDR block:2407:c080:1200:217e::/64 An EIP is required if you want to access DB instances through a public network. View EIP 🔀 Addresses available: 251 View In-use IP Addresses Database Port Default port: 3306 Security Group (?) Q View Security Group [2] Ensure that port 3306 of the security group allows traffic from your server IP address to the DB instance. Create Security Group

Figure 5-7 Configuring network information as planned





Step 3 Confirm the settings.

- To modify your settings, click **Previous**.
- If there is no need to modify your settings, click Submit.

Step 4 Bind an EIP to the created instance.

1. On the **Instances** page, click the instance name to go to the **Overview** page.

Figure 5-9 Locating your instance in the list



- 2. In the navigation pane on the left, choose **Connectivity & Security**. In the **Connection Information** area, click **Bind** next to the **EIP** field.
- 3. In the displayed dialog box, all unbound EIPs are listed. Select the EIP you have created for the instance and click **OK**.

Figure 5-10 Binding an EIP



----End

5.1.5 Configuring an RDS for MySQL Instance in the DR Center

5.1.5.1 Creating a VPC and Security Group

Create a VPC and security group for the DR instance to be configured, ensuring that it is in a different region from the instance created for production center.

Creating a VPC

- **Step 1** Go to the **Create VPC** page.
- **Step 2** On that page, select **AP-Singapore** for **Region**, and configure the basic information, subnet, and IP address.

Figure 5-11 Creating a VPC


Step 3 Click Create Now.

----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 **AP-Singapore**.
- **Step 3** Under the service list, choose **Networking** > **Virtual Private Cloud**.
- **Step 4** In the navigation pane on the left, choose **Access Control** > **Security Groups**.
- **Step 5** Click **Create Security Group**.

Figure 5-12 Creating a security group



Step 6 Click Create Now.

----End

5.1.5.2 Creating an RDS for MySQL Instance

Create an RDS for MySQL instance as a DR instance and select the VPC you configured for the instance.

Procedure

- Step 1 Go to the Buy DB Instance page.
- **Step 2** Select **AP-Singapore** for **Region**. Configure instance information and click **Buy**.

Figure 5-13 Selecting a DB engine





Figure 5-14 Selecting specifications






Figure 5-15 Configuring network information as planned




Figure 5-16 Additional options

Step 3 Confirm the settings.

- To modify your settings, click **Previous**.
- If there is no need to modify your settings, click **Submit**.

----End

5.1.6 Configuring Remote Disaster Recovery

5.1.6.1 Creating a DRS Disaster Recovery Task

Create a DRS disaster recovery task in the same region as the RDS for MySQL instance configured for the DR center.

Procedure

- **Step 1** Go to the **Create Disaster Recovery Task** page.
- Step 2 Select AP-Singapore for Region. Set Disaster Recovery Relationship to Current cloud as standby, and DR DB Instance to the RDS for MySQL DR instance created in the AP-Singapore region, and click Create Now.




Figure 5-17 Setting DR instance information

Step 3 Return to the **Disaster Recovery Management** page and check the status of the task.

----End

5.1.6.2 Configuring the Disaster Recovery Task

Configure the disaster recovery task, including setting the source and destination databases.

Procedure

- **Step 1** On the **Disaster Recovery Management** page, locate the created disaster recovery task and click **Edit** in the **Operation** column.
- **Step 2** Add the EIP of the DRS instance to the inbound rule of the security group associated with the RDS for MySQL instance in the production center, select TCP, and set the port number to that of the RDS for MySQL instance of the production center.


Figure 5-18 Adding a security group rule



In the **Source Database** area, set **IP Address or Domain Name** and **Port** to the EIP and port of the RDS for MySQL instance in the production center. When the connection test is successful, click **Next**.


Figure 5-19 Editing a disaster recovery task

Source Database



Step 3 Configure the flow control and click **Next**.

Figure 5-20 Configuring flow control




- **Step 4** Check the disaster recovery task. When the check success rate reaches 100%, click **Next**.
- **Step 5** Configure parameters and click **Next**.

Figure 5-21 Configuring parameters



Step 6 Configure Start Time and click Submit.

Figure 5-22 Starting the task



Step 7 On the **Disaster Recovery Management** page, check the task status. The status is **Disaster recovery in progress**.

For a task that is in the **Disaster recovery in progress** state, you can use **data comparison** to check whether data is consistent before and after the disaster recovery.

----End


5.1.6.3 Performing a Primary/Standby Switchover

If the source database in the production center is faulty, manually switch the DR instance to the read/write state. Then, data is written to the DR instance and synchronized to the source database.

Procedure

- **Step 1** Find that the source database in the production center is faulty. For example, the source database cannot be connected, the source database execution is slow, or the CPU usage is high.
- **Step 2** Receive an SMN email notification.
- **Step 3** Check the delay of the DR task.
- **Step 4** Check that the services of the source database have been stopped. For details, see **How Do I Ensure that All Services on the Database Are Stopped?**
- **Step 5** Select the task, click the **Batch Operations** drop-down box in the upper left corner and select **Primary/Standby Switchover**.

Figure 5-23 Primary/standby switchover



Step 6 Change the database IP address on your application and use it to connect to the database. Then data is properly read from and written to the database.

----End

6 Security Best Practices

Security is a shared responsibility between Huawei Cloud and you. Huawei Cloud is responsible for the security of cloud services to provide a secure cloud. As a tenant, you should properly use the security capabilities provided by cloud services to protect data, and securely use the cloud. For details, see **Shared Responsibilities**.

This section provides actionable guidance for enhancing the overall security of using DRS for data migration. You can continuously evaluate the security status of your DRS tasks for secure data migration and enhance their overall security defense. By doing this, only the minimum permissions required for business can be assigned, and data can be protected from leakage and tampering both in transmission.

Make security configurations from the following dimensions to meet your business needs.

- Using Fine-Grained Authorization to Control the Usage Scope of DRS Resources
- Using Secure and Reliable Networks and Encrypted Transmission Protocols
- Using Network Access Control to Isolate the Network for Data Synchronization
- Configuring Independent Database Migration Users and Assigning the Minimum Permissions
- Creating HA Tasks to Improve Service Availability
- Properly Using Authentication Credentials to Prevent Data Leaks

Using Fine-Grained Authorization to Control the Usage Scope of DRS Resources

1. Set only the minimum permissions for IAM users with different roles to prevent data leakage or misoperations caused by excessive permissions.

To better isolate and manage permissions, you are advised to configure an independent IAM administrator and grant them the permission to manage IAM policies. The IAM administrator can create different user groups based on your service requirements. User groups correspond to different data access

scenarios. By adding users to user groups and binding IAM policies to user groups, the IAM administrator can grant different data access permissions to employees in different departments based on the principle of least privilege.

2. Fine-grained authorization is recommended to enable fine-grained control on user permissions.

Fine-grained policies define permissions by APIs. You are advised to **create a custom policy** based on your DRS operation permissions.

Using Secure and Reliable Networks and Encrypted Transmission Protocols

1. You are advised to use a secure network, such as a VPN, for data synchronization.

Do not use an EIP network if possible. Instead, use a secure network, such as a VPN, for data transmission. Configure firewalls, security groups, and ACL rules to reduce the attack surface and improve the network security for data synchronization.

2. The certificate+SSL connection mode is recommended.

The certificate+SSL mode is a secure connection mode. It protects the integrity and confidentiality of data during transmission, but slightly affects the read and write performance of the database. In certain scenarios that are sensitive to synchronization performance, you need to balance performance and security.

Using Network Access Control to Isolate the Network for Data Synchronization

Firewalls, Access Control List (ACL) rules, and security groups are used for network access control to effectively control the network range for DRS to access databases and isolate the network for data synchronization from other networks, ensuring DRS task security.

Configuring Independent Database Migration Users and Assigning the Minimum Permissions

If you migrate data as user **root** or other service users, permission control may be disordered and permission leakage may occur. When creating a DRS task, you are advised to create independent migration accounts for the source and destination databases and grant the minimum permissions to the migration accounts by referring to the user guide to reduce the risk of account and permission leaks.

Creating HA Tasks to Improve Service Availability

DRS provides the cross-AZ HA. If the instance in the primary AZ becomes faulty, services can be switched over to the instance in the standby AZ to continue data replication.

Properly Using Authentication Credentials to Prevent Data Leaks

When you use code or API Explorer to call APIs, you need to obtain a token using the account password or AK/SK information. You need to comply with the secure

encoding rules, properly manage authentication credentials, and do not hardcode authentication information in plaintext.