

Data Replication Service

Getting Started

Issue 01
Date 2024-12-26



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

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

1 Creating a Real-Time Migration Task.....	1
2 Creating a Real-Time Synchronization Task.....	15
3 Creating a Real-Time DR Task.....	30
4 Creating a Workload Replay Task.....	44
5 Common Tasks.....	56

1 Creating a Real-Time Migration Task

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

This section describes how to migrate full+incremental data from an on-premises MySQL database to a Huawei Cloud RDS for MySQL instance through a public network. For more information about real-time migration, see [Migration Overview](#).

Video Tutorial

Operation Process

Process	Description
Preparations	Sign up for a HUAWEI ID, complete real-name authentication, enable Huawei Cloud services, and prepare the database to be migrated, the database connection account, and a network.
Step 1: Create a Migration Task	Select the source and destination databases, and create and start a migration task.
Step 2: Query Migration Progress	Check the migration progress.
Step 3: Compare Migration Items	Create a comparison task to check whether the data in the source database is consistent with that in the destination database.
Stopping a Task	After confirming that the data is consistent, cut over workloads and stop the DRS task.

Preparations

Before creating a real-time migration task, prepare the Huawei account, database to be migrated, database connection account, and network environment by referring to the following steps.

Registering a HUAWEI ID

Prepare a Huawei account, create a user, and grant permissions to the user to use DRS.

Step 1 Access [Huawei Cloud website](#).

Step 2 Click **Sign Up** and follow the instructions to create your account (your HUAWEI ID).

Step 3 Select the service agreement and click **Enable**.

The system displays a message indicating that Huawei Cloud services have been enabled.

Step 4 Perform real-name authentication.

- Individual account: [Individual Real-Name Authentication](#)
- Enterprise account: [Enterprise Real-Name Authentication](#)

Step 5 For fine-grained permissions management, create an Identity and Access Management (IAM) user and user group on the IAM console and grant the user specific operation permissions. For details, see [Creating a User and Granting Permissions](#)

----End

Databases

Before creating a real-time migration task, you need to prepare the source and destination databases.

- The source database in this example is an on-premises MySQL database. Prepare the following source database details:

Item	Example Value	Description
DB engine version	MySQL 5.7	-
IP address	10.154.217.42	Enter an IP address.
Port	3306	-

- The destination database in this example is a Huawei Cloud RDS for MySQL instance. You need to prepare the database details by referring to [Getting Started with RDS for MySQL](#).

Item	Example Value	Description
Region	AP-Singapore	To reduce network latency, select the region nearest to you.
Instance name	rds-mysql	Specify a name that will be easy to identify.
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	-

Connection Accounts

To ensure a smooth data migration, you are advised to create an independent database account for connecting to the DRS task. In addition, you need to ensure that the account is granted required permissions.

- Source database user
 - a. Log in to the source database.
 - b. Run the following statement to create database user **usersrc**:
CREATE USER 'usersrc'@'%' IDENTIFIED BY 'password';
 - c. Run the following statements to grant permissions to **usersrc**:
GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* TO 'usersrc';
GRANT SELECT, SHOW VIEW, EVENT, LOCK TABLES ON *.* TO 'usersrc';
- Destination database user
 - a. [Connect to an RDS for MySQL instance through DAS](#).
 - b. Run the following statement to create database user **usertar**. For details, see [Creating a Database Account Through DAS](#).
CREATE USER 'usertar'@'%' IDENTIFIED BY 'password';
 - c. Run the following statement to grant permissions to **usertar**:

GRANT SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER ON *.* TO 'usertar' WITH GRANT OPTION;

Network

Before creating a migration task over a public network, you need to apply for an EIP and configure the firewall of your local data center so that the EIP can access the on-premises database.

Step 1 Create an EIP for the DRS migration task by referring to [Assigning an EIP](#).

Step 2 Configure the firewall of the local data center.

Add the EIP created in [Step 1](#) to the whitelist of the source database to ensure that the source database can access EIP.

Step 3 Configure the IP address whitelist for the on-premises database.

Add the DRS instance EIP to the whitelist of the on-premises database to allow the access from the DRS instance.

----End

Step 1: Create a Migration Task

This step describes how to set parameters based on the example values in [Preparations](#). For more information about real-time migration, see [Migration Overview](#).

Step 1 Go to the [Create Migration Task](#) page.

Step 2 Configure basic information as follows:

Figure 1-1 Basic information

Region: AP-Singapore

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

Project: AP-Singapore

* Task Name: DRS-Migration

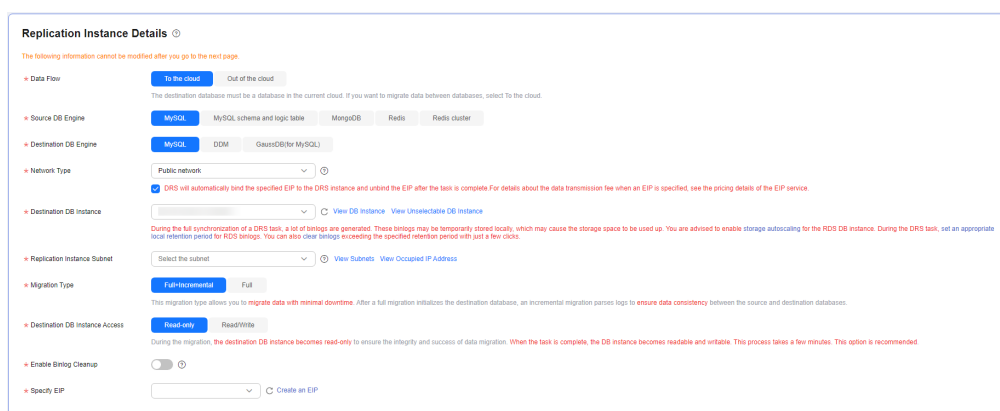
Description: [Empty text area]

Parameter	Example Value	Description
Region	AP-Singapore	The region where the current task is deployed.
Project	AP-Singapore	The project corresponds to the current region.

Parameter	Example Value	Description
Task Name	DRS-Migration	The name of the migration task.
Description	Leave this parameter blank for now.	Task description.

Step 3 Configure information about the replication instance.

Figure 1-2 Replication instance details



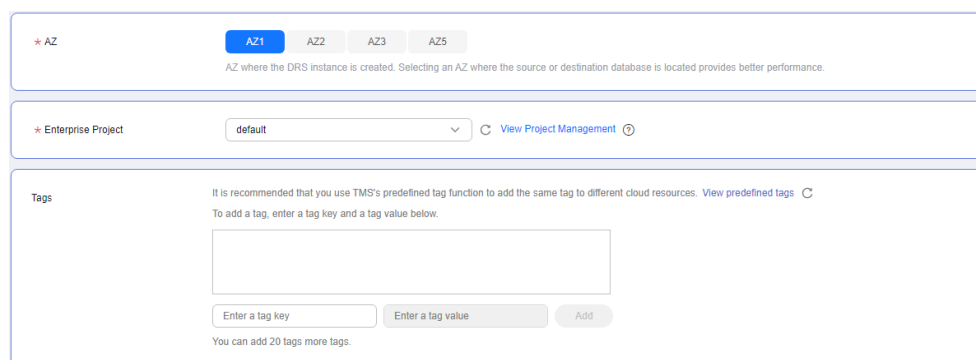
Parameter	Example Value	Description
Data Flow	To the Cloud	The direction of the migration task. The options are To the cloud and Out of the cloud . <ul style="list-style-type: none"> To the cloud: The destination database is a Huawei Cloud DB instance and data needs to be transferred to the cloud. Out of the cloud: The source database is a Huawei Cloud DB instance and data needs to be transferred out of the cloud.
Source DB Engine	MySQL	The DB engine type of the source database.
Destination DB Engine	MySQL	The DB engine type of the destination database.

Parameter	Example Value	Description
Network Type	Public network	The network type of the migration task. <ul style="list-style-type: none">• VPC: suitable for migrations between cloud databases of the same account in the same region and VPC.• Public network: suitable for migrations from on-premises databases or external cloud databases to destination databases.• VPN or Direct Connect: suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
Destination DB Instance	rds-mysql	Select a destination DB instance. The destination DB instance is rds-mysql in Databases .
Replication Instance Subnet	Default subnet	The subnet where the migration task is.
Migration Type	Full +Incremental	The available options are Full and Full +Incremental . <ul style="list-style-type: none">• Full: This migration type is suitable for scenarios where services can be interrupted. If you are performing a full migration, do not perform operations on the source database during the migration, or 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.
Destination DB Instance Access	Read-only	During the migration, the destination DB instance can be set to read-only or read/write. <ul style="list-style-type: none">• Read-only: During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and improves the success rate of data migration.• Read/Write: During the migration, the destination database can be read and written.

Parameter	Example Value	Description
Enable Binlog Cleanup	No	Whether to enable the function of quickly clearing binlogs of the destination database.
Specify EIP	-	Select an EIP for the DRS instance. In this example, it is the EIP created in Network .

Step 4 Select an AZ and an enterprise project.

Figure 1-3 AZ and Enterprise Project



Parameter	Example Value	Description
AZ	AZ1	Select the AZ where you want to create the DRS task.
Enterprise Project	default	If your account is associated with an enterprise project, select the project from the Enterprise Project drop-down list. For more details, see Enterprise Management User Guide .
Tags	Leave this parameter blank for now.	Tags for the migration task. Adding tags helps you better identify and manage your tasks.

Step 5 Click **Create Now**.

Wait for 5 to 10 minutes. If the message "Replication instance created successfully" is displayed, the DRS task is successfully created.

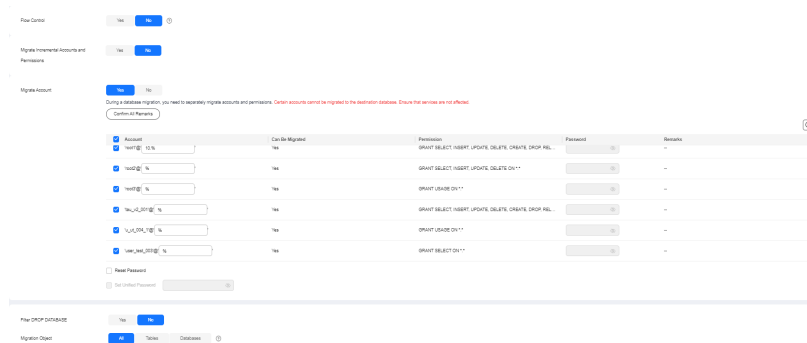
Step 6 On the **Configure Source and Destination Databases** page, specify source and destination database details and click **Test Connection** for both the source and destination databases to confirm they have been connected to the DRS instance.

1. In the **Source Database** area, enter the source database details in [Connection Accounts](#).

2. Click **Test Connection**.
If the message "Test successful" is displayed, the source database is connected.
3. In the **Destination Database** area, enter the destination database details in **Connection Accounts**.
4. Click **Test Connection**.
If the message "Test successful" is displayed, the destination database is connected.
5. Click **Next**.

Step 7 On the **Set Task** page, select the objects to be migrated and click **Next**.

Figure 1-4 Migration Type



Parameter	Example Value	Description
Flow Control	No	Whether to set a maximum migration speed of the DRS task. This function is disabled by default. Flow Control is only used during the full migration phase.
Migrate Incremental Accounts and Permissions	No	Whether to migrate incremental accounts and permissions during database migration.
Migrate Account	No	Whether to migrate users in the source database.
Filter DROP DATABASE	No	To reduce the risks involved in a migration, DRS allows you to filter out DROP operations.
Migration Object	All	Select objects to be migrated. You can select All , Tables , or Databases as required.

- Step 8** On the **Check Task** page, check the migration task.
- If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.

- If any of the checks fail, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

If the check success rate reaches 100%, the pre-check is successful.

Step 9 Click **Next**

Step 10 On the **Compare Parameters** page, compare and modify common and performance parameters.

- If you do not want to compare parameters or the parameters of the source database are the same as those of the destination database, skip this step.
- If the parameters of the source database are inconsistent with those of the destination database, you can manually change the values of the destination database parameters or click **Use Source Database Value** to make the parameter values of the source and destination database be the same.

Step 11 Click **Next**

Step 12 On the **Confirm Task** page, specify **Start Time**.

Figure 1-5 Task startup settings

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

★ Send Notifications ?

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

Parameter	Example Value	Description
Start Time	Start upon task creation	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.
Send Notifications	Disable	This parameter is optional. After enabled, select an SMN topic. If the task billing is about to start, the status, latency metric, or data of the migration task is abnormal, DRS will send you a notification.
Stop Abnormal Tasks After	14	Any task in the Abnormal state that has run for longer than the period you set here (in days) will automatically stop.

Step 13 Confirm that the configured information is correct and click **Submit** to submit the task.

Step 14 After the migration task is submitted, go to the **Online Migration Management** page to check the task status.

- If the task status is **Starting**, the task has been started.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you reconfigure the task, DRS applies for resources for the task again.

----End

Step 2: Query Migration Progress

After the migration task is started, you can check the migration progress. DRS shows the migration progress using a progress bar, so you can track the migration progress in real time.

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

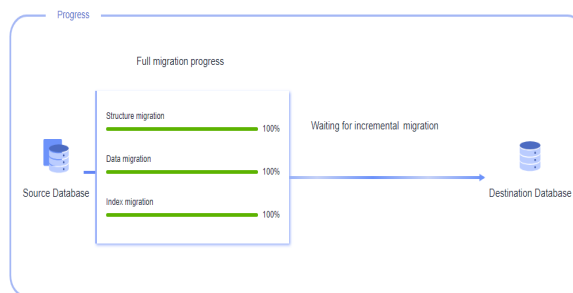
The **Basic Information** page is displayed.

Step 2 Choose **Migration Progress** to check the progress and details.

- **Progress Overview:** When a full migration is complete, the progress of each item reaches 100%. The migration progress is displayed based on the number of migrated objects.

Figure 1-6 Migration progress

Start the task and get 7 days free. If the status indicates a synchronization or migration is complete, there may still be triggers or events to be migrated before the entire task is finished.
Last Updated May 30, 2024 16:05:54 GMT+08:00 Source Position:mysql-bin.033658:13675 Consumed Position:NA



- **Migration Details:** If the values of **Total Items** and **Migrated Items** are the same, the object migration is complete.

Figure 1-7 Migration progress details

Migration Details

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

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

- **Incremental migration delay:** After the full migration is complete, you can check the delay of the incremental migration on the **Migration Progress** page.
Delay refers to the delay from when the transaction was submitted to the source database to when it is synchronized to the destination database and executed.
If the delay is 0, the source database is consistent with the destination database, and no new transactions need to be synchronized.

----End

Step 3: Compare Migration Items

For a full+incremental migration task, when the full migration is complete and the incremental migration delay is 0, you can create a comparison task to check whether the data in the source database is consistent with that in the destination database.

- **Object-level comparison:** It compares databases, indexes, tables, views, stored procedures and functions, and sorting rules of tables.
- **Data-level comparison:** It checks the consistency of rows or values in the migration tables.

Object-Level Comparison

You can create an object-level comparison task to check the integrity of database objects.

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

The **Basic Information** page is displayed.


Step 2 Choose **Migration Comparison**.

The **Migration Comparison** page is displayed.

Step 3 Click **Object-Level Comparison**.

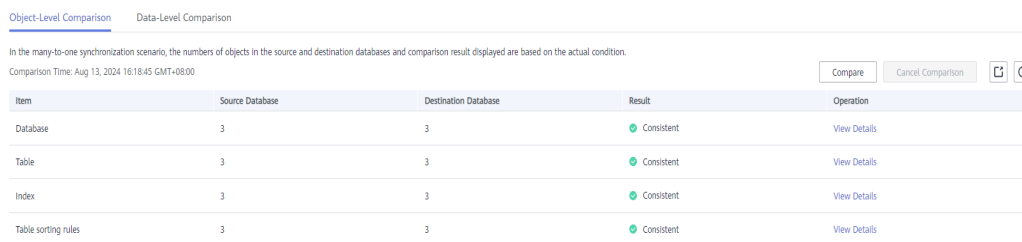
The **Object-Level Comparison** page is displayed.

Step 4 Click **Compare**.

Step 5 After 5 to 10 minutes, click  to check the comparison result of each comparison item.

If **Consistent** is displayed in the **Result** column, the object migration is complete.

Figure 1-8 Object-Level Comparison



Object-Level Comparison Data-Level Comparison

In the many-to-one synchronization scenario, the numbers of objects in the source and destination databases and comparison result displayed are based on the actual condition.
Comparison Time: Aug 13, 2024 16:18:45 GMT+08:00

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

----End

Data-Level Comparison

After the database object comparison is complete, you can create a data-level comparison task to check the number of rows or values of the migrated data.

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

The **Basic Information** page is displayed.

Step 2 Choose **Migration Comparison**.

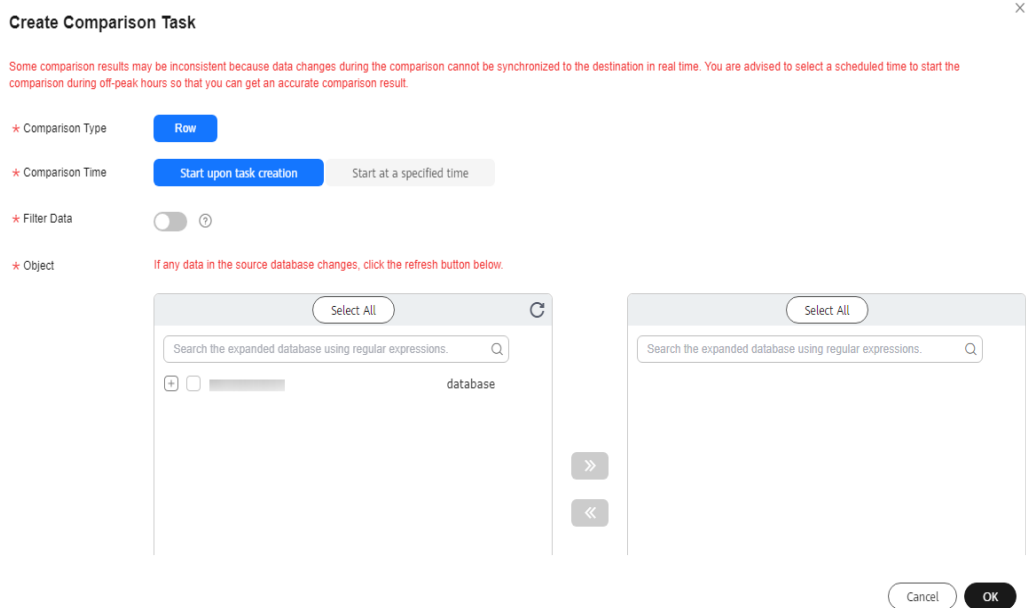
The **Migration Comparison** page is displayed.

Step 3 In the **Before You Start** pane, click **Validate All Rows/Values**.

The **Create Comparison Task** page is displayed.


Step 4 Select a comparison policy and click **OK** to submit the comparison task.

Figure 1-9 Creating a comparison task




Parameter	Example Value	Description
Comparison Type	Row	The available options are Row and Value . <ul style="list-style-type: none"> Row comparison: It checks whether the number of rows in a source table is the same as that in the migrated table. Value comparison: It checks whether the data in a source table is consistent with that in the migrated table.
Comparison Time	Start upon task creation	The available options are Start upon task creation and Start at a specified time .
Filter Data	Disable	Whether to set filtering criteria for comparison objects. After this function is enabled, objects can be compared based on the configured filtering criteria.
Object	Select All	Select objects to be compared as needed.

Step 5 Go back to the **Data-Level Comparison** tab.

Step 6 Click  to refresh the list. After the comparison task is complete, you can check the comparison result.

If **Consistent** is displayed in the **Result** column, the object data in the source database is consistent with that in the destination database.

Figure 1-10 Row comparison result



----End

Stopping a Task

After confirming that the source and destination data is consistent, you can cut over workloads to the destination database. After workload cutover is successfully completed and applications are running properly, you can stop the migration task.

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

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

----End

2 Creating a Real-Time Synchronization Task

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

This section describes how to synchronize full+incremental data from an on-premises MySQL database to a Huawei Cloud RDS for MySQL instance through a public network. For more information about real-time synchronization, see [Synchronization Overview](#).

Video Tutorial

Operation Process

Process	Description
Preparations	Sign up for a HUAWEI ID, complete real-name authentication, enable Huawei Cloud services, and prepare the database to be synchronized, the database connection account, and a network.
Step 1: Create a Synchronization Task	Select the source and destination databases, and create and start a synchronization task.
Step 2: Query Synchronization Progress	Check the synchronization progress.
Step 3: Compare Synchronization Items	Create a comparison task to check whether the data in the source database is consistent with that in the destination database.
Stopping a Task	After confirming that the data is consistent, cut over workloads and stop the DRS task.

Preparations

Before creating a real-time synchronization task, prepare the Huawei account, database to be synchronized, database connection account, and network environment by referring to the following steps.

Registering a HUAWEI ID

Prepare a Huawei account, create a user, and grant permissions to the user to use DRS.

Step 1 Access [Huawei Cloud website](#).

Step 2 Click **Sign Up** and follow the instructions to create your account (your HUAWEI ID).

Step 3 Select the service agreement and click **Enable**.

The system displays a message indicating that Huawei Cloud services have been enabled.

Step 4 Perform real-name authentication.

- Individual account: [Individual Real-Name Authentication](#)
- Enterprise account: [Enterprise Real-Name Authentication](#)

Step 5 For fine-grained permissions management, create an Identity and Access Management (IAM) user and user group on the IAM console and grant the user specific operation permissions. For details, see [Creating a User and Granting Permissions](#)

----End

Databases

Before creating a real-time synchronization task, you need to prepare the source and destination databases.

- The source database in this example is an on-premises MySQL database. Prepare the following source database details:

Item	Example Value	Description
DB engine version	MySQL 5.7	-
IP address	10.154.217.42	Enter an IP address.
Port	3306	-

- The destination database in this example is a Huawei Cloud RDS for MySQL instance. You need to prepare the database details by referring to [Getting Started with RDS for MySQL](#).

Item	Example Value	Description
Region	AP-Singapore	To reduce network latency, select the region nearest to you.
Instance name	rds-mysql	Specify a name that will be easy to identify.
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	-

Connection Accounts

To ensure a smooth data synchronization, you are advised to create an independent database account for connecting to the DRS task. In addition, you need to ensure that the account is granted required permissions.

- Source database user
 - a. Log in to the source database.
 - b. Run the following statement to create database user **usersrc**:
CREATE USER 'usersrc'@'%' IDENTIFIED BY 'password';
 - c. Run the following statements to grant permissions to **usersrc**:
GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* TO 'usersrc';
GRANT SELECT, SHOW VIEW, EVENT, LOCK TABLES, ON [Database to be synchronized].* TO 'usersrc';
- Destination database user
 - a. [Connect to an RDS for MySQL instance through DAS.](#)
 - b. Run the following statement to create database user **usertar**. For details, see [Creating a Database Account Through DAS.](#)
CREATE USER 'usertar'@'%' IDENTIFIED BY 'password';
 - c. Run the following statement to grant permissions to **usertar**:

```
GRANT SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER,  
REFERENCES ON *.* TO 'usertar';
```

Network

Before creating a synchronization task over a public network, you need to apply for an EIP and configure the firewall of your local data center so that the EIP can access the on-premises database.

Step 1 Create an EIP for the DRS synchronization task by referring to [Assigning an EIP](#).

Step 2 Configure the firewall of the local data center.

Add the EIP created in [Step 1](#) to the whitelist of the source database to ensure that the source database can access EIP.

Step 3 Configure the IP address whitelist for the on-premises database.

Add the DRS instance EIP to the whitelist of the on-premises database to allow the access from the DRS instance.

----End

Step 1: Create a Synchronization Task

This step describes how to set parameters based on the example values in [Preparations](#). For more information about real-time synchronization, see [Synchronization Overview](#).

Step 1 Go to the [Create Synchronization Task](#) page.

Step 2 Configure basic information as follows:

Figure 2-1 Basic information

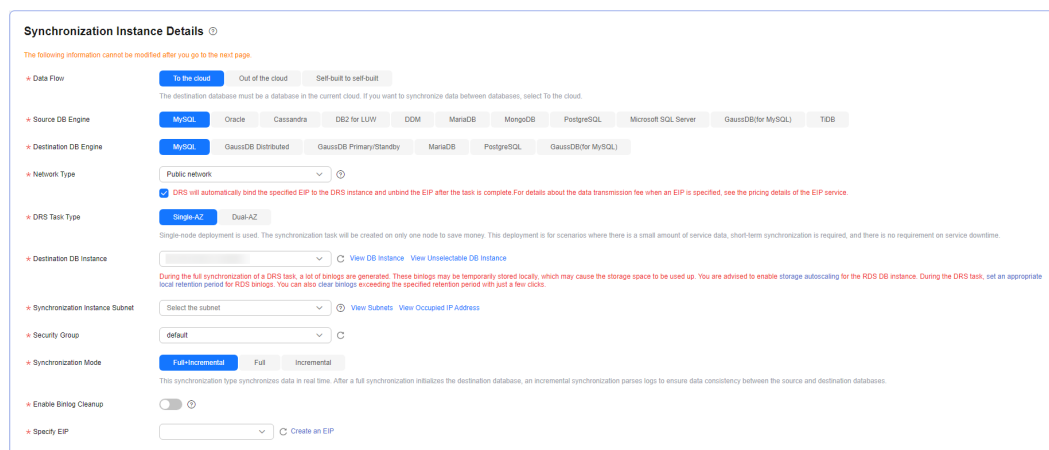
The screenshot shows a configuration form with the following fields and options:

- Billing Mode:** Radio buttons for 'Yearly/Monthly' and 'Pay-per-use' (selected).
- Region:** A dropdown menu showing 'AP-Singapore' with a location pin icon. Below it, a note reads: 'Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.'
- Project:** A dropdown menu showing 'AP-Singapore'.
- * Task Name:** A text input field containing 'DRS-Synchronization' and a help icon.
- Description:** A large text area that is currently empty, with a character count '0/256' and a help icon at the bottom right.

Parameter	Example Value	Description
Billing Mode	Pay-per-use	DRS billing mode. <ul style="list-style-type: none"> ● Yearly/Monthly: A prepaid billing mode in which you pay for resources before using it. Bills are settled based on the subscription period. ● Pay-per-use billing: A postpaid billing mode. In this mode, you pay for what you used.
Region	AP-Singapore	The region where the current task is deployed.
Project	AP-Singapore	The project corresponds to the current region.
Task Name	DRS-Synchronization	The name of the synchronization task.
Description	Leave this parameter blank for now.	Task description.

Step 3 Configure information about the synchronization instance.

Figure 2-2 Synchronization instance details



Parameter	Example Value	Description
Data Flow	To the Cloud	The direction of the synchronization task. The options are To the cloud , Out of the cloud , and Self-built to self-built . <ul style="list-style-type: none">• To the cloud: The destination database is a Huawei Cloud DB instance and data needs to be transferred to the cloud.• Out of the cloud: The source database is a Huawei Cloud DB instance and data needs to be transferred out of the cloud.• Self-built to self-built: Neither the source database nor the destination databases is a Huawei Cloud DB instance.
Source DB Engine	MySQL	The DB engine type of the source database.
Destination DB Engine	MySQL	The DB engine type of the destination database.
Network Type	Public network	The network type of the synchronization task. <ul style="list-style-type: none">• VPC: suitable for synchronizations between cloud databases of the same account in the same region and VPC.• Public network: suitable for synchronizations from on-premises databases or external cloud databases to destination databases.• VPN or Direct Connect: suitable for synchronizations from on-premises databases to cloud databases or between cloud databases across regions using a VPN, Direct Connect, Cloud Connect, VPCEP, or a VPC peering connection.
DRS Task Type	Single-AZ	Type of the DRS task. The value can be Single-AZ or Dual-AZ .
Destination DB Instance	rds-mysql	Select a destination DB instance. The destination DB instance is rds-mysql in Databases .
Synchronization Instance Subnet	Default subnet	The subnet where the synchronization task is.
Synchronization Mode	Full +Incremental	The available options are Full , Full +Incremental , and Incremental .
Enable Binlog Cleanup	No	Whether to enable the function of quickly clearing binlogs of the destination database.

Parameter	Example Value	Description
Specify EIP	-	Select an EIP for the DRS instance. In this example, it is the EIP created in Network .

Step 4 Select the task specifications and AZ.

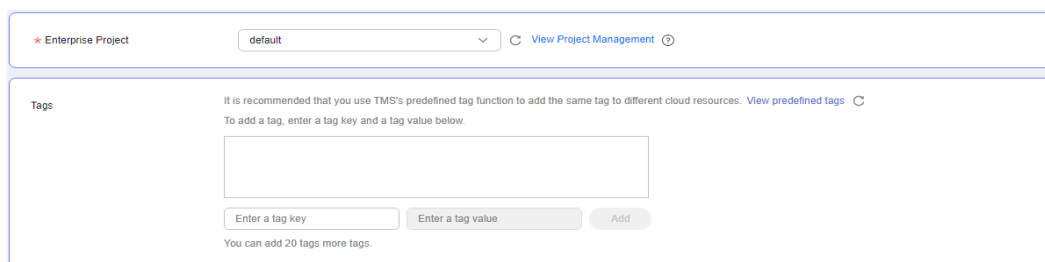
Figure 2-3 Specifications and AZ



Parameter	Example Value	Description
Specifications	Large	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time Synchronization .
AZ	AZ1	Select the AZ where you want to create the DRS task.

Step 5 Configure an enterprise project and tags.

Figure 2-4 Enterprise Project and Tags



Parameter	Example Value	Description
Enterprise Project	default	If your account is associated with an enterprise project, select the project from the Enterprise Project drop-down list. For more details, see Enterprise Management User Guide .

Parameter	Example Value	Description
Tags	Leave this parameter blank for now.	Tags for the synchronization task. Adding tags helps you better identify and manage your tasks.

Step 6 Click **Create Now**.

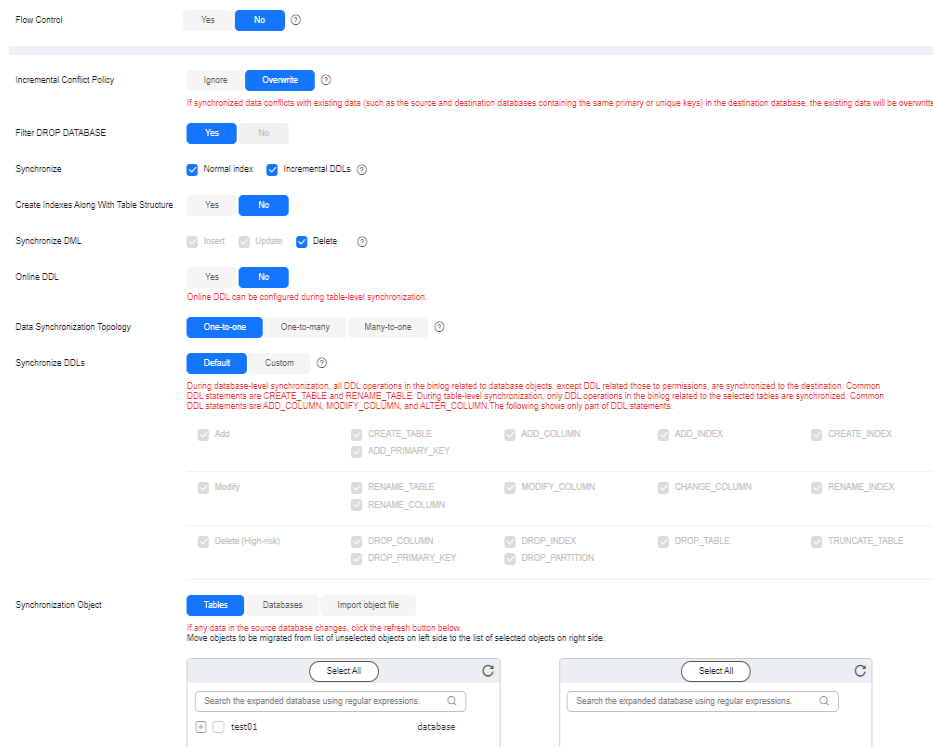
Wait for 5 to 10 minutes. If the message "Synchronization instance created successfully" is displayed, the DRS task is successfully created.

Step 7 On the **Configure Source and Destination Databases** page, specify source and destination database details and click **Test Connection** for both the source and destination databases to confirm they have been connected to the DRS instance.

1. In the **Source Database** area, enter the source database details in [Connection Accounts](#).
2. Click **Test Connection**.
If the message "Test successful" is displayed, the source database is connected.
3. In the **Destination Database** area, enter the destination database details in [Connection Accounts](#).
4. Click **Test Connection**.
If the message "Test successful" is displayed, the destination database is connected.
5. Click **Next**.

Step 8 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Figure 2-5 Synchronization Mode



Parameter	Example Value	Description
Flow Control	No	Whether to set a maximum synchronization speed of the DRS task. This function is disabled by default. Flow Control is only used during the full synchronization phase.
Incremental Conflict Policy	Overwrite	The policy for handling data conflicts during incremental synchronization. By default, conflicts in the full synchronization phase are ignored.
Filter DROP DATABASE	No	To reduce the risks involved in a synchronization, DRS allows you to filter out DROP operations.
Synchronize	Common index and Incremental DDLs	Select the type of the objects to be synchronized.
Create Indexes Along With Table Structure	No	Whether to create indexes along with the table structure in the full synchronization phase.

Parameter	Example Value	Description
Synchronize DML	Insert, Update, and Delete	Select the DML operations to be synchronized. By default, all DML operations are selected.
Online DDL	No	If table-level synchronization is selected, you can choose whether to synchronize Online DDL. By default, Online DDL is not synchronized.
Data Synchronization Topology	One-to-one	Data synchronization supports multiple synchronization topologies. You can plan your synchronization instances based on service requirements. For details, see Data Synchronization Topologies .
Synchronize DDLs	Default	Select DDL type for incremental synchronization.
Synchronization Object	Tables	Select objects to be synchronized. You can select All , Tables , or Databases as required.

Step 9 On the **Process Data** page, set the filtering rules for data processing.

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

Step 10 On the **Check Task** page, check the synchronization task.

- If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.
- If any of the checks fail, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

If the check success rate reaches 100%, the pre-check is successful.

Step 11 Click **Next**

Step 12 On the **Confirm Task** page, specify **Start Time**.

Figure 2-6 Task startup settings

The screenshot shows three settings:

- Start Time:** A radio button is selected for "Start upon task creation" (highlighted in blue), and "Start at a specified time" is unselected. A help icon (?) is visible next to the second option.
- Send Notifications:** A toggle switch is turned off. A help icon (?) is visible next to it.
- Stop Abnormal Tasks After:** A text input field contains the value "14". A help icon (?) is visible next to it, followed by the text: "Abnormal tasks run longer than the period you set (unit: day) will automatically stop."

Parameter	Example Value	Description
Start Time	Start upon task creation	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.
Send Notifications	Disable	This parameter is optional. After enabled, select an SMN topic. If the status, latency metric, or data of the synchronization task is abnormal, DRS will send you a notification.
Stop Abnormal Tasks After	14	Any task in the Abnormal state that has run for longer than the period you set here (in days) will automatically stop.

Step 13 Confirm that the configured information is correct and click **Submit** to submit the task.

Step 14 After the synchronization task is submitted, go to the **Data Synchronization Management** page to check the task status.

- If the task status is **Starting**, the task has been started.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you reconfigure the task, DRS applies for resources for the task again.

----End

Step 2: Query Synchronization Progress

After the synchronization task is started, you can check the synchronization progress. DRS shows the synchronization progress using a progress bar, so you can track the synchronization progress in real time.


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

The **Basic Information** page is displayed.

Step 2 Choose **Synchronization Progress** to check the progress and details.

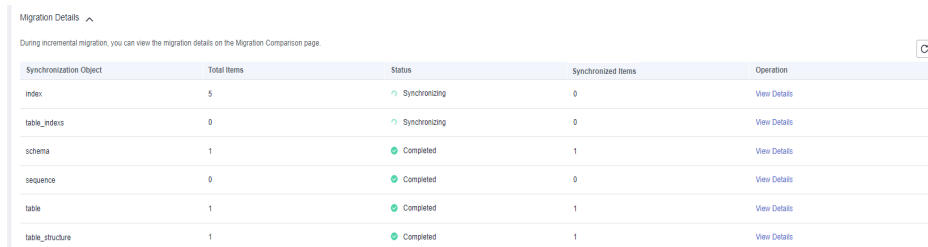
- **Progress:** When the synchronization progress reaches 100%, full synchronization is complete. The synchronization progress is displayed based on the number of synchronized objects.

Figure 2-7 Synchronization progress



- **Migration Details:** If the values of **Total Items** and **Synchronized Items** are the same, the object synchronization is complete.

Figure 2-8 Synchronization progress details



Migration Details

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

Synchronization Object	Total Items	Status	Synchronized Items	Operation
index	5	Synchronizing	0	View Details
table_indexes	0	Synchronizing	0	View Details
schema	1	Completed	1	View Details
sequence	0	Completed	0	View Details
table	1	Completed	1	View Details
table_structure	1	Completed	1	View Details

- **Incremental synchronization delay:** After the full synchronization is complete, you can check the delay of the incremental synchronization on the **Synchronization Progress** page.

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

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

----End

Step 3: Compare Synchronization Items

For a full+incremental synchronization task, when the full synchronization is complete and the incremental synchronization delay is 0, you can create a comparison task to check whether the data in the source database is consistent with that in the destination database.

- **Object-level comparison:** It compares databases, indexes, tables, views, stored procedures and functions, and sorting rules of tables.
- **Data-level comparison:** It checks the consistency of rows or values in the synchronization tables.

Object-Level Comparison

You can create an object-level comparison task to check the integrity of database objects.

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

The **Basic Information** page is displayed.


Step 2 Choose **Synchronization Comparison**.

The **Synchronization Comparison** page is displayed.

Step 3 Click **Object-Level Comparison**.

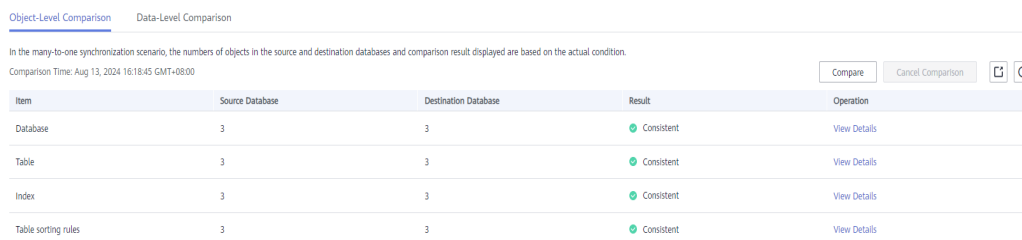
The **Object-Level Comparison** page is displayed.

Step 4 Click **Compare**.





Step 5 After 5 to 10 minutes, click  to check the comparison result of each comparison item.

If **Consistent** is displayed in the **Result** column, the object synchronization is complete.

Figure 2-9 Object-Level Comparison



The screenshot shows a web interface for 'Object-Level Comparison'. At the top, there are two tabs: 'Object-Level Comparison' (selected) and 'Data-Level Comparison'. Below the tabs, a note states: 'In the many-to-one synchronization scenario, the numbers of objects in the source and destination databases and comparison result displayed are based on the actual condition.' Below this note, the 'Comparison Time' is 'Aug 13, 2024 16:18:45 GMT+08:00'. On the right side, there are buttons for 'Compare', 'Cancel Comparison', and a refresh icon. The main content is a table with the following data:

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

----End

Data-Level Comparison

After the database object comparison is complete, you can create a data-level comparison task to check the number of rows or values of the synchronized data.

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

The **Basic Information** page is displayed.

Step 2 Choose **Synchronization Comparison**.

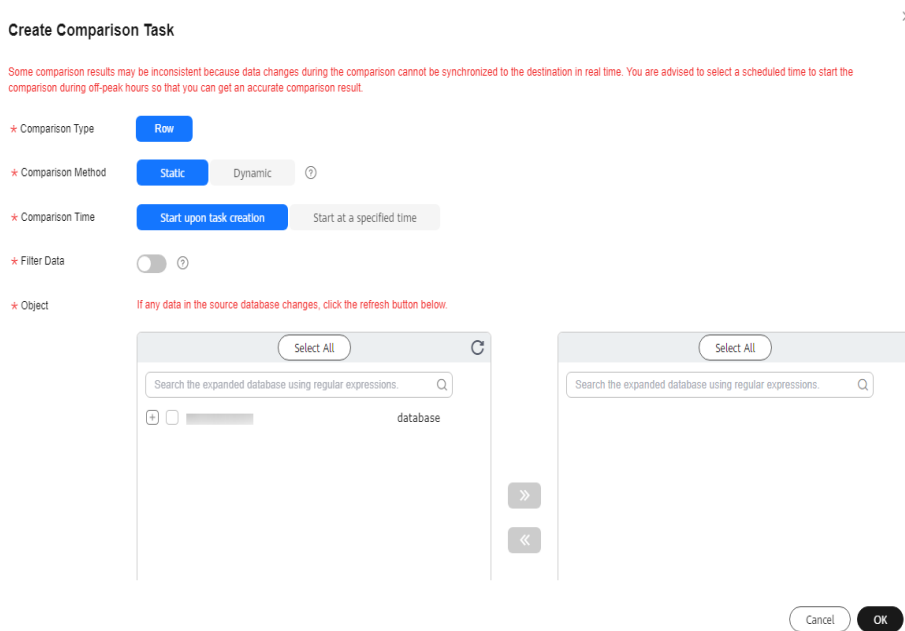
The **Synchronization Comparison** page is displayed.

Step 3 Click the **Data-Level Comparison** tab.

Step 4 Click **Create Comparison Task**.

Step 5 Select a comparison policy and click **OK** to submit the comparison task.


Figure 2-10 Creating a comparison task



Parameter	Example Value	Description
Comparison Type	Row	The available options are Row and Value . <ul style="list-style-type: none"> • Row comparison: It checks whether the number of rows in a source table is the same as that in the synchronized table. • Value comparison: It checks whether the data in a source table is consistent with that in the synchronized table.
Comparison Policy	One-to-one	DRS supports one-to-one and many-to-one comparison policies. <ul style="list-style-type: none"> • One-to-one: compares the number of rows in a table in the source database with that in the table mapped to the destination database. • Many-to-one: compares the number of rows in a table in the source database with that in the aggregate table mapped to the destination database.
Comparison Time	Start upon task creation	The available options are Start upon task creation and Start at a specified time .
Filter Data	Disable	Whether to set filtering criteria for comparison objects. After this function is enabled, objects can be compared based on the configured filtering criteria.

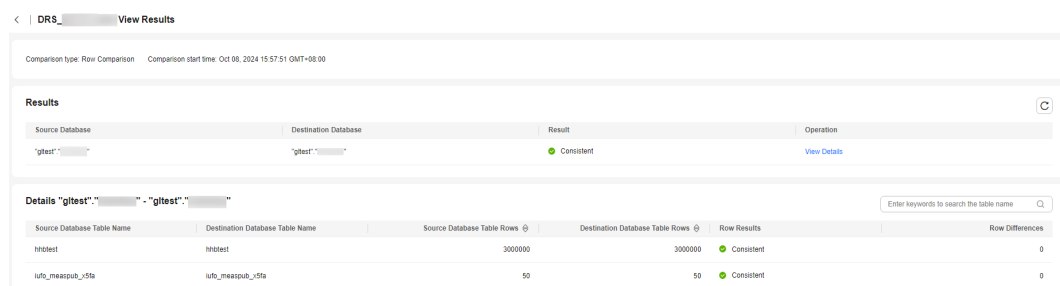
Parameter	Example Value	Description
Object	Select All	Select objects to be compared as needed.

Step 6 Go back to the **Data-Level Comparison** tab.

Step 7 Click  to refresh the list. After the comparison task is complete, you can check the comparison result.

If **Consistent** is displayed in the **Result** column, the object data in the source database is consistent with that in the destination database.

Figure 2-11 Row comparison result



----End

Stopping a Task

After confirming that all data in the source database is synchronized to the destination database, you can stop the synchronization task.

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

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

----End

3 Creating a Real-Time DR Task

To prevent service interruptions caused by regional faults, DRS provides disaster recovery to ensure service continuity. If the region where the primary instance is located encounters a natural disaster and cannot be connected, you can switch the remote DR instance to the primary instance. To reconnect to the primary instance, you only need to change the connection address on the application side. DRS allows you to perform cross-region real-time synchronization between a primary DB instance and a DR instance during disaster recovery.

This section uses two RDS for MySQL DB instances in different regions as an example to describe how to quickly create a remote single-active DR task through the public network.

Operation Process

Process	Description
Preparations	Sign up for a HUAWEI ID, complete real-name authentication, enable Huawei Cloud services, and prepare the database to be migrated, the database connection account, and a network.
Step 1: Create a DR Task	Select the source and destination databases as required and create a DR task.
Step 2: Query DR Progress	During the disaster recovery, check the DR progress.
Step 3: Compare DR Items	Compare objects and data to be synchronized to ensure data consistency.
Step 4 (Optional): Perform a DR Switchover	Perform a primary/standby switchover for the DR task.

Preparations

Before creating a real-time DR task, prepare the Huawei account, database to be migrated, database connection account, and network environment by referring to the following steps.

Registering a HUAWEI ID

Prepare a Huawei account, create a user, and grant permissions to the user to use DRS.

Step 1 Access [Huawei Cloud website](#).

Step 2 Click **Sign Up** and follow the instructions to create your account (your HUAWEI ID).

Step 3 Select the service agreement and click **Enable**.

The system displays a message indicating that Huawei Cloud services have been enabled.

Step 4 Perform real-name authentication.

- Individual account: [Individual Real-Name Authentication](#)
- Enterprise account: [Enterprise Real-Name Authentication](#)

Step 5 For fine-grained permissions management, create an Identity and Access Management (IAM) user and user group on the IAM console and grant the user specific operation permissions. For details, see [Creating a User and Granting Permissions](#)

----End

Databases

Before creating a real-time DR task, you need to prepare the source and destination databases.

- The source database in this example is an RDS for MySQL instance in the CN-Hong Kong region. For details, see [Buy an RDS for MySQL DB Instance](#).

Item	Example Value	Description
Region	CN-Hong Kong	To reduce network latency, select the region nearest to you.
Instance name	rds-DRsrc	Specify a name that will be easy to identify.
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.

Item	Example Value	Description
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	-
EIP	10.154.217.42	Enter an IP address.

- The destination database in this example is an RDS for MySQL instance in the AP-Singapore region. For details, see [Buy an RDS for MySQL DB Instance](#).

Item	Example Value	Description
Region	AP-Singapore	To reduce network latency, select the region nearest to you.
Instance name	rds-DRtar	Specify a name that will be easy to identify.
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	-

Connection Accounts

To ensure a smooth data disaster recovery, you are advised to create an independent database account for connecting to the DRS task. In addition, you need to ensure that the account is granted required permissions.

- Source database user
 - a. [Connect to an RDS for MySQL instance through DAS.](#)
 - b. Run the following statement to create database user **usersrc**. For details, see [Creating a Database Account Through DAS.](#)
CREATE USER 'usersrc'@'%' IDENTIFIED BY 'password';
 - c. Run the following statements to grant permissions to **usersrc**:
GRANT
 SELECT,CREATE,ALTER,DROP,DELETE,INSERT,UPDATE,TRIGGER,REFERENC
 ES,SHOW VIEW,EVENT,INDEX,LOCK TABLES,CREATE VIEW,CREATE
 ROUTINE,ALTER ROUTINE,CREATE USER,RELOAD,REPLICATION
 SLAVE,REPLICATION CLIENT ON *.* TO 'usersrc'@'%' **WITH GRANT
 OPTION;**
- Destination database user
 - a. [Connect to an RDS for MySQL instance through DAS.](#)
 - b. Run the following statement to create database user **usertar**. For details, see [Creating a Database Account Through DAS.](#)
CREATE USER 'usertar'@'%' IDENTIFIED BY 'password';
 - c. Run the following statement to grant permissions to **usertar**:
GRANT
 SELECT,CREATE,ALTER,DROP,DELETE,INSERT,UPDATE,TRIGGER,REFERENC
 ES,SHOW VIEW,EVENT,INDEX,LOCK TABLES,CREATE VIEW,CREATE
 ROUTINE,ALTER ROUTINE,CREATE USER,RELOAD,REPLICATION
 SLAVE,REPLICATION CLIENT ON *.* TO 'usertar'@'%' **WITH GRANT
 OPTION;**


Network

Before creating a DR task over a public work, you need to apply for an EIP and set security group rules for the source database so that the EIP can access the source database.

Step 1 Create an EIP for the DR task by referring to [Assigning an EIP](#).

Step 2 Configure security group rules for the source database.

Add the EIP in [Step 1](#) of the DRS instance to the inbound rule of the security group associated with the [source RDS for MySQL database](#), select TCP, and set the port number to that of the source database.



----End

Step 1: Create a DR Task

This step describes how to set parameters based on the example values in [Preparations](#). For more information about real-time disaster recovery, see [DR Overview](#).

Step 1 Go to the [Create Disaster Recovery Task](#) page.

Step 2 Configure basic information as follows:

Figure 3-1 Basic information

The screenshot shows a configuration form with the following fields and values:

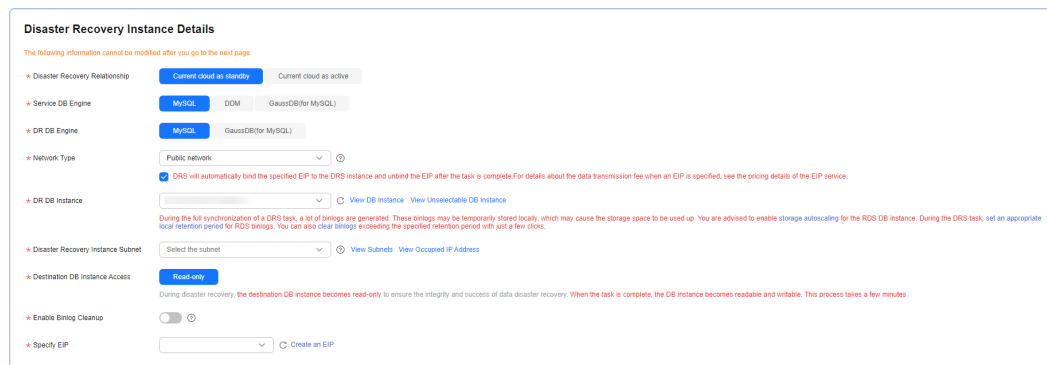
- Billing Mode:** Radio buttons for 'Yearly/Monthly' and 'Pay-per-use' (selected).
- Region:** A dropdown menu showing 'AP-Singapore'.
- Project:** A dropdown menu showing 'AP-Singapore'.
- Task Name:** A text input field containing 'DRS-DRTask'.
- Description:** A large text area that is currently empty.

Below the Region dropdown, there is a note: "Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region."

Parameter	Example Value	Description
Billing Mode	Pay-per-use	DRS billing mode. <ul style="list-style-type: none"> ● Yearly/Monthly: A prepaid billing mode in which you pay for resources before using it. Bills are settled based on the subscription period. ● Pay-per-use billing: A postpaid billing mode. In this mode, you pay for what you used.
Region	AP-Singapore	The region where the current task is deployed.
Project	AP-Singapore	The project corresponds to the current region.
Task Name	DRS-DRTask	The name of the DR task.
Description	Leave this parameter blank for now.	Task description.

Step 3 Configure information about the disaster recovery instance.

Figure 3-2 Disaster recovery instance details




Parameter	Example Value	Description
Disaster Recovery Relationship	Current cloud as standby	The disaster recovery relationship of the DR task. The options are Current cloud as standby and Current cloud as active . <ul style="list-style-type: none"> Current cloud as standby: The DR database is a Huawei Cloud DB instance and data needs to be transferred to the cloud. Current cloud as active: The service database is a Huawei Cloud DB instance and data needs to be transferred out of the cloud.
Service DB Engine	MySQL	The DB engine type of the service database.
DR DB Engine	MySQL	The DB engine type of the DR database.
Network Type	Public network	The network type of the DR task.
DRS Task Type	Single-AZ	Type of the DRS task. The value can be Single-AZ or Dual-AZ .
DR DB Instance	rds-DRtar	Select a DR DB instance. The DR DB instance is rds-DRtar in Databases .
Disaster Recovery Instance Subnet	Default subnet	The subnet where the DR task is.
Destination DB Instance Access	Read-only	During the disaster recovery, set the DR DB instance to the read-only state.
Enable Binlog Cleanup	No	Whether to enable the function of quickly clearing binlogs of the DR database.

Parameter	Example Value	Description
Specify EIP	-	Select an EIP for the DRS instance. In this example, it is the EIP created in Network .

Step 4 Select the task specifications and AZ.


Figure 3-3 Specifications and AZ



Parameter	Example Value	Description
Specifications	Large	DRS instance specifications. Different specifications have different performance upper limits. For details, see Real-Time DR .
AZ	AZ1	Select the AZ where you want to create the DRS task.

Step 5 Configure an enterprise project and tags.

Figure 3-4 Enterprise Project and Tags



Parameter	Example Value	Description
Enterprise Project	default	If your account is associated with an enterprise project, select the project from the Enterprise Project drop-down list. For more details, see Enterprise Management User Guide .

Parameter	Example Value	Description
Tags	Leave this parameter blank for now.	Tags for the DR task. Adding tags helps you better identify and manage your tasks.

Step 6 Click **Create Now**.

Wait for 5 to 10 minutes. If the message "Disaster recovery instance created successfully" is displayed, the DRS task is successfully created.


Step 7 On the **Configure Source and Destination Databases** page, specify source and destination database details and click **Test Connection** for both the source and destination databases to confirm they have been connected to the DRS instance.

1. In the **Source Database** area, enter the source database details in [Connection Accounts](#).
2. Click **Test Connection**.
If the message "Test successful" is displayed, the source database is connected.
3. In the **Destination Database** area, enter the destination database details in [Connection Accounts](#).
4. Click **Test Connection**.
If the message "Test successful" is displayed, the destination database is connected.
5. Click **Next**.

Step 8 On the **Configure DR** page, specify flow control and click **Next**.

- Set **Flow Control** to **No**.
- Set **Migrate Definer to User** to **Yes**.

Figure 3-5 DR settings



Parameter	Example Value	Description
Flow Control	No	Whether to set a maximum speed for the DR task. This function is disabled by default.

Parameter	Example Value	Description
Migrate Definer to User	Yes	Whether to migrate the Definers of all source database objects to the destination database user entered during the connection test.

Step 9 On the **Check Task** page, check the DR task.

- If there are any items that require confirmation, view and confirm the details first before proceeding to the next step.
- If any of the checks fail, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.

If the check success rate reaches 100%, the pre-check is successful.

Step 10 Click **Next**

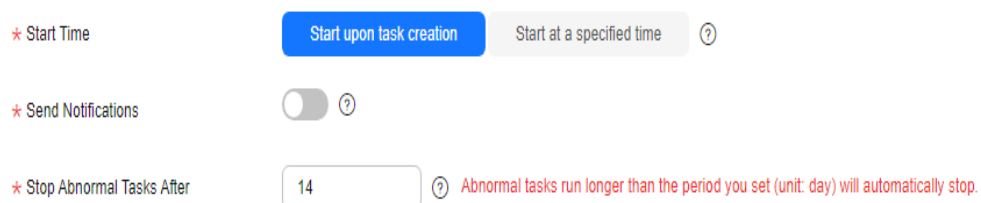
Step 11 On the **Compare Parameters** page, compare and modify common and performance parameters.

- If you do not want to compare parameters or the parameters of the source database are the same as those of the destination database, skip this step.
- If the parameters of the source database are inconsistent with those of the destination database, you can manually change the values of the destination database parameters or click **Use Source Database Value** to make the parameter values of the source and destination database be the same.

Step 12 Click **Next**

Step 13 On the **Confirm Task** page, specify **Start Time**.

Figure 3-6 Task startup settings



Parameter	Example Value	Description
Start Time	Start upon task creation	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.

Parameter	Example Value	Description
Send Notifications	Disable	This parameter is optional. After enabled, select an SMN topic. If the status, latency metric, or data of the DR task is abnormal, DRS will send you a notification.
Stop Abnormal Tasks After	14	Any task in the Abnormal state that has run for longer than the period you set here (in days) will automatically stop.

Step 14 Confirm that the configured information is correct and click **Submit** to submit the task.

Step 15 After the DR task is submitted, go to the **Disaster Recovery Management** page to check the task status.

- If the task status is **Starting**, the task has been started.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you reconfigure the task, DRS applies for resources for the task again.

----End

Step 2: Query DR Progress

After the DR task is started, you can check the DR progress. DRS shows the DR progress using a progress bar, so you can track the DR progress in real time.

Step 1 On the **Disaster Recovery Management** page, click the target DR task in the **Task Name/ID** column.

The **Basic Information** page is displayed.

Step 2 Choose **Disaster Recovery Progress** to check the progress.

- When the data initialization is complete, the initialization progress is displayed as 100%.
- Delay refers to the delay from when the transaction was submitted to the source database to when it is synchronized to the destination database and executed.

When the delay is 0, data is synchronized from the service database to the DR database in real time. You can view more metrics, such as RPO and RTO, on the **Disaster Recovery Monitoring** tab.

Step 3 When the delay of the DR task is 0s, you can use **data comparison** to check whether the data in the service database is consistent with that in the DR database.

----End

Step 3: Compare DR Items

When the task enters the **Disaster recovery in progress** state and the RPO and RTO are 0, you can create a comparison task to check whether the data in the source database is consistent with that in the destination database.

- Object-level comparison: It compares databases, indexes, tables, views, stored procedures and functions, and sorting rules of tables. You are advised to perform an object-level comparison first.
- Data-level comparison: It checks the consistency of rows or values in the migration tables.

Object-Level Comparison

You can create an object-level comparison task to check the integrity of database objects.

Step 1 On the **Disaster Recovery Management** page, click the target DR task in the **Task Name/ID** column.

The **Basic Information** page is displayed.


Step 2 Choose **Disaster Recovery Comparison**.

The **Disaster Recovery Comparison** page is displayed.

Step 3 Click **Object-Level Comparison**.

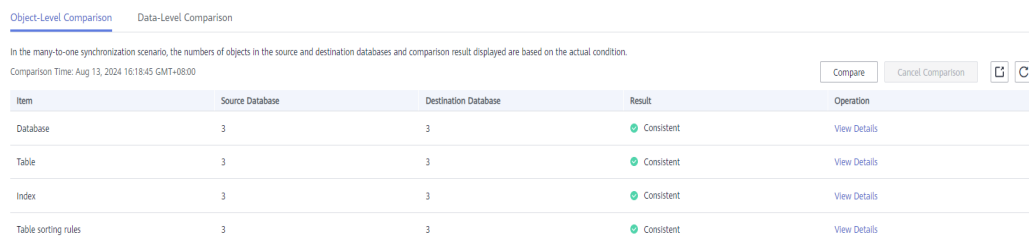
The **Object-Level Comparison** page is displayed.

Step 4 Click **Compare**.

Step 5 After 5 to 10 minutes, click  to check the comparison result of each comparison item.

If **Consistent** is displayed in the **Result** column, the object migration is complete.

Figure 3-7 Object-Level Comparison



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

----End

Data-Level Comparison

After the database object comparison is complete, you can create a data-level comparison task to check the number of rows or values of the disaster recovery data.

Step 1 On the **Disaster Recovery Management** page, click the target DR task in the **Task Name/ID** column.

The **Basic Information** page is displayed.

Step 2 Choose **Disaster Recovery Comparison**.

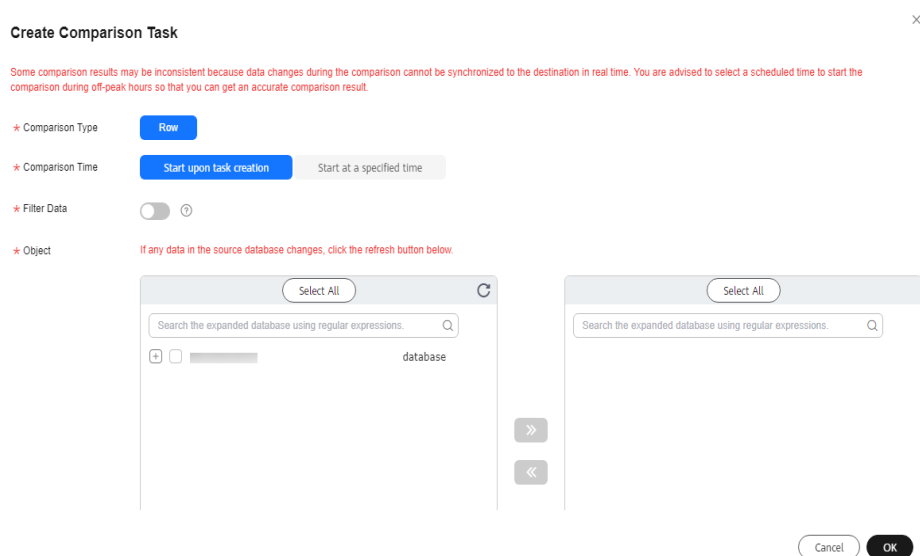
The **Disaster Recovery Comparison** page is displayed.

Step 3 Click the **Data-Level Comparison** tab.

Step 4 Click **Create Comparison Task**.

Step 5 In the displayed dialog box, select a comparison policy and click **OK**.


Figure 3-8 Creating a comparison task



Parameter	Example Value	Description
Comparison Type	Row	The available options are Row and Value . <ul style="list-style-type: none"> Row comparison: It checks whether the number of rows in a service table is the same as that in the DR table. Value comparison: It checks whether the data in a service table is consistent with that in the DR table.
Comparison Time	Start upon task creation	The available options are Start upon task creation and Start at a specified time .
Filter Data	Disable	Whether to set filtering criteria for comparison objects. After this function is enabled, objects can be compared based on the configured filtering criteria.


Parameter	Example Value	Description
Object	Select All	Select objects to be compared as needed.

Step 6 Go back to the **Data-Level Comparison** tab.

Step 7 Click  to refresh the list. After the comparison task is complete, you can check the comparison result.

If **Consistent** is displayed in the **Result** column, the object data in the source database is consistent with that in the destination database.

Figure 3-9 Row comparison result



----End

Step 4 (Optional): Perform a DR Switchover

DRS allows you to perform a switchover for a DR task. When a disaster occurs, the DR database can be promoted to the service database to ensure service continuity.

- Before a switchover, services are running properly in the service database and data is synchronized to the DR database in real time. In this case, data cannot be written into the DR database.
- After a switchover, the DR database becomes readable and writable, services can be switched to the DR database, and data cannot be written to the service database.

Step 1 On the **Disaster Recovery Management** page, locate the target DR task.

Step 2 Click the task name.

The **Basic Information** page is displayed.

Step 3 Choose **Disaster Recovery Progress** to check the RPO and RTO.

- RPO measures the difference between the data in the service database and the data in the DRS instance. When RPO is 0, all the data in the service database has been migrated to the DRS instance.
- RTO measures the amount of data being transmitted. When RTO is 0, all transactions on the DRS instance have been completed on the DR database.

If both RPO and RTO are 0, data has been completely migrated to the DR database. Then, you can determine whether to perform a switchover.

Step 4 Choose **Disaster Recovery Monitoring**.

Step 5 A switchover can be performed only when the task status is **Disaster recovery in progress**.

- Click **Promote Current Cloud** to promote the current instance to the service database.
- Click **Demote Current Cloud** to demote the current instance to the DR database.

The DR relationship involves only one primary database. During a primary/secondary switchover, ensure that there is no data written to the database that will be the standby node, and no data will be written to the standby node in the future. The data of the standby node is synchronized only from the primary node. Any other write operations will pollute the data in the standby database, data conflicts occur in the database and cannot be resolved.

----End

4 Creating a Workload Replay Task

A workload replay task simulates the service load of the source database on the destination database so you can evaluate the effectiveness and performance of the destination database. It is typically used in the following scenarios:

- **Function testing:** By creating a workload replay task, you can evaluate how the service load of the source database runs on the destination database before database migration.
- **Peak load testing:** By specifying the replay thread and speed, you can simulate the peak service load of the source database and analyze the stability of the destination database when workloads increase sharply.

This section uses two RDS for MySQL DB instances in different regions as an example to describe how to replay SQL workload files over the public network.

Video Tutorial

Operation Process

Process	Description
Preparations	Sign up for a HUAWEI ID, complete real-name authentication, enable Huawei Cloud services, and prepare the workload files and the destination database, the database connection account, and a network.
Step 1: Create a Workload Replay Task	Select the source and destination databases as required and create a workload replay task.
Step 2: Query Replay Progress	During the workload replay, query the progress.
Step 3: Check Replay Reporting	After workload replay is complete, check the replay reporting to learn about the execution of each SQL statement.
Stopping a Task	After confirming that the workload replay task is no longer used, stop the DRS task.

Preparations

Before creating a workload replay task, prepare the Huawei account, source and destination databases, database connection accounts, SQL workload files, and network environment by referring to the following steps.

Registering a HUAWEI ID

Prepare a Huawei account, create a user, and grant permissions to the user to use DRS.

Step 1 Access [Huawei Cloud website](#).

Step 2 Click **Sign Up** and follow the instructions to create your account (your HUAWEI ID).

Step 3 Select the service agreement and click **Enable**.

The system displays a message indicating that Huawei Cloud services have been enabled.

Step 4 Perform real-name authentication.

- Individual account: [Individual Real-Name Authentication](#)
- Enterprise account: [Enterprise Real-Name Authentication](#)

Step 5 For fine-grained permissions management, create an Identity and Access Management (IAM) user and user group on the IAM console and grant the user specific operation permissions. For details, see [Creating a User and Granting Permissions](#)

----End

Databases

Before creating a workload replay task, you need to prepare the source and destination databases.

- The source database in this example is an RDS for MySQL instance in the CN-Hong Kong region. For details, see [Buy an RDS for MySQL DB Instance](#).

Item	Example Value	Description
Region	CN-Hong Kong	To reduce network latency, select the region nearest to you.
Instance name	rds-Replaysrc	Specify a name that will be easy to identify.
DB engine version	MySQL 5.7	-

Item	Example Value	Description
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	-
EIP	10.154.217.42	Enter an IP address.


- The destination database in this example is an RDS for MySQL instance in the AP-Singapore region. For details, see [Buy an RDS for MySQL DB Instance](#).

Item	Example Value	Description
Region	AP-Singapore	To reduce network latency, select the region nearest to you.
Instance name	rds-Replaytar	Specify a name that will be easy to identify.
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	-

SQL Workload Files

Before creating a task, you need to capture SQL workload files of 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 **CN-Hong Kong**.

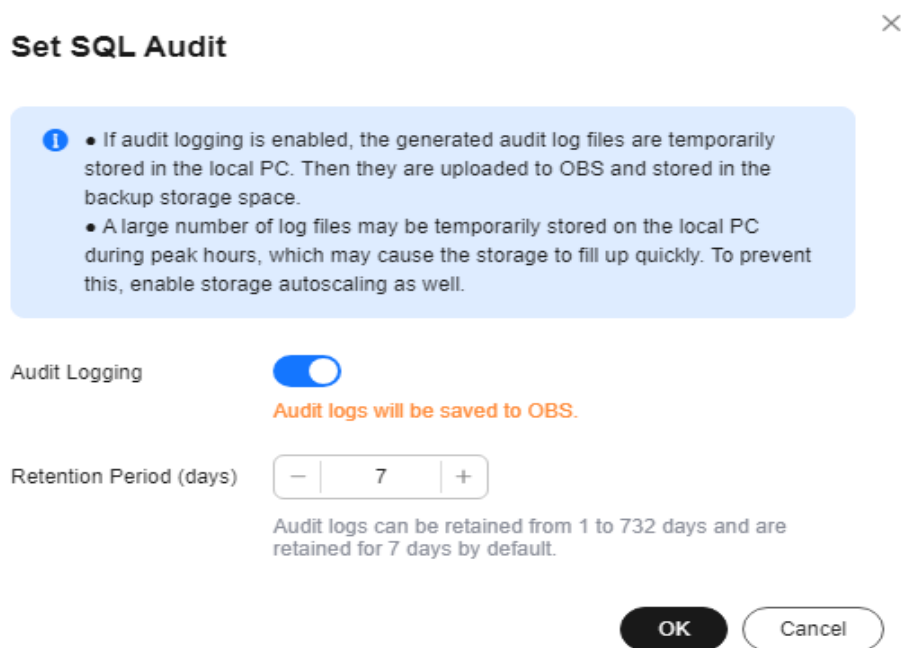
Step 3 Click **Service List**. Under **Databases**, click **Relational Database Service**. The RDS console is displayed.

Step 4 On the **Instances** page, click the name of the source DB instance created in [Databases](#) to go to the **Overview** page.

Step 5 In the navigation pane, choose **SQL Audits**. On the displayed page, click **Set SQL Audit** above the list. In the displayed dialog box, configure information as required and click **OK**.

- **Audit Logging:** Enable
- **Retention Period (days):** 7

Figure 4-1 Set SQL Audit



----End

Connection Accounts

To ensure a smooth workload replay, you are advised to create an independent database account for connecting to the DRS task. It is recommended that the permissions of the destination database user be the same as that of the source database user.

Destination database user

Step 1 [Connect to an RDS for MySQL instance through DAS.](#)

Step 2 Run the following statement to create database user **usertar**. For details, see [Creating a Database Account Through DAS.](#)

```
CREATE USER 'usertar'@'%' IDENTIFIED BY 'password';
```

Step 3 Run the following statement to grant permissions to **usertar**:

```
GRANT
SELECT,CREATE,ALTER,DROP,DELETE,INSERT,UPDATE,TRIGGER,REFERENCES,SHOW
VIEW,EVENT,INDEX,LOCK TABLES,CREATE VIEW,CREATE ROUTINE,ALTER
ROUTINE,CREATE USER,RELOAD,REPLICATION SLAVE,REPLICATION CLIENT ON *.*
TO 'usertar'@'%' WITH GRANT OPTION;
```

----End

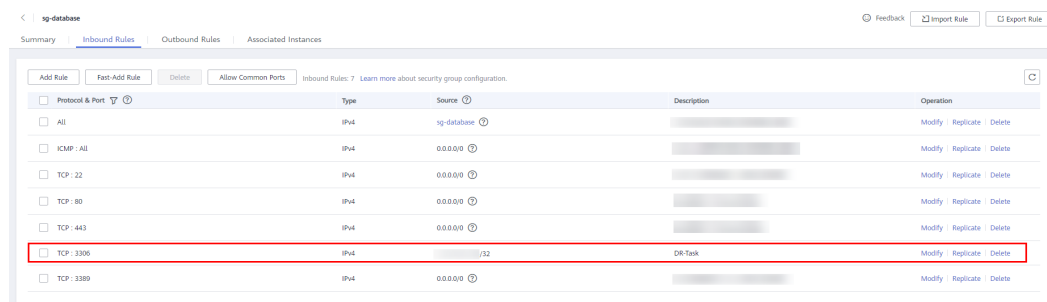
Network

Before creating a workload replay task over a public work, you need to apply for an EIP and set security group rules for the source database so that the EIP can access the source database.

Step 1 Create an EIP for the workload replay task by referring to [Assigning an EIP.](#)

Step 2 Configure security group rules for the source database.

Add the EIP in [Step 1](#) of the DRS instance to the inbound rule of the security group associated with the [source RDS for MySQL database](#), select TCP, and set the port number to that of the source database.



----End

Step 1: Create a Workload Replay Task

This step describes how to set parameters based on the example values in [Preparations](#). For more information about workload replay, see [Replay Overview](#).

Step 1 Go to the [Create Workload Replay Task](#) page.

Step 2 Configure basic information as follows:

Figure 4-2 Basic information

Region

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

Project

* Task Name

Description

0/256 ↕

Parameter	Example Value	Description
Region	AP-Singapore	The region where the current task is deployed.
Project	AP-Singapore	The project corresponds to the current region.
Task Name	DRS-Replay	The name of the workload replay task.
Description	Leave this parameter blank for now.	Task description.

Step 3 Configure information about the replay instance.

Figure 4-3 Replay instance details

Replay Instance Details

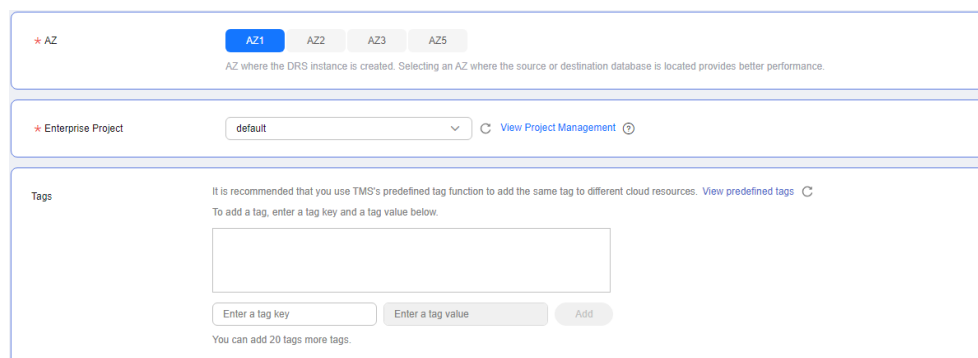
The following information cannot be modified after you go to the next page.

- Data Flow: Current cloud To the cloud
- Source DB Engine: MySQL GaussDB(for MySQL)
- Workload File Source: Download from Huawei Cloud AP's Download from Huawei Cloud OBS
- Destination DB Engine: MySQL GaussDB(for MySQL)
- Network Type:
- DRS will automatically bind the specified EIP to the DRS instance and unbind the EIP after the task is complete. For details about the data transmission fee when an EIP is specified, see the pricing details of the EIP service.
- Destination DB Instance: [View DB Instance](#) [View Unselectable DB Instance](#)
- Replay instance Subnet: [View Subnets](#) [View Occupied IP Address](#)
- Specify EIP:

Parameter	Example Value	Description
Data Flow	Current cloud	The direction of the workload replay task. The options are Current cloud and To the cloud . <ul style="list-style-type: none"> • Current cloud: Both the source and destination databases are Huawei Cloud DB instances. • To the cloud: The destination database is a Huawei Cloud DB instance and data needs to be transferred to the cloud.
Source DB Engine	MySQL	The DB engine type of the source database.
Workload File Source	Download from Huawei Cloud APIs	Specifies where the workload file in the source database is from.
Destination DB Engine	MySQL	The DB engine type of the destination database.
Network Type	Public network	The network type of the workload replay task.
Destination DB Instance	rds-Replaytar	Select a destination DB instance. The destination DB instance is rds-Replaytar in Databases .
Replay Instance Subnet	Default subnet	The subnet where the workload replay task is.
Specify EIP	-	Select an EIP for the DRS instance. In this example, it is the EIP created in Network .

Step 4 Select an AZ and an enterprise project.

Figure 4-4 AZ and Enterprise Project



Parameter	Example Value	Description
AZ	AZ1	Select the AZ where you want to create the DRS task.
Enterprise Project	default	If your account is associated with an enterprise project, select the project from the Enterprise Project drop-down list. For more details, see Enterprise Management User Guide .
Tags	Leave this parameter blank for now.	Tags for the workload replay task. Adding tags helps you better identify and manage your tasks.

Step 5 Click **Create Now**.

Wait for 5 to 10 minutes. If the message "Replay instance created successfully" is displayed, the DRS task is successfully created.

Step 6 On the **Configure Source and Destination Databases** page, specify the source and destination database details.

- Source database settings
 - **Workload File Source: Download from Huawei Cloud APIs**
 - **DB Instance Name:** Select the source DB instance **rds-Replaysrc** created in [Databases](#).
 - **Workload Type: Audit log**
 - **Time Range:** Select the time range for audit logs in [SQL Workload Files](#).
- Destination database settings
Enter the connection details about the destination database in [Connection Accounts](#).

Step 7 Configure the task.

Figure 4-5 Task settings

Task Settings

SQL Type

Replay Mode Performance Transaction ?

Filter out SQLs ?
Add
 You can add 9 more SQLs.

Filter out SQLs Without Conditions ?

Maximum Concurrent Connections ?

Acceleration Configuration ?

Parameter	Example Value	Description
SQL Type	SELECT	Select the SQL type to be replayed to the destination database. The default value is SELECT . The available options are SELECT , INSERT , UPDATE , DELETE , and DDL .
Replay Mode	Performance	Mode of the workload replay task. You can select Performance or Transaction . <ul style="list-style-type: none"> In performance mode, you can set how many concurrent connections are allowed. SQL statements are replayed to the destination database based on a set number of connections. The SQL execution sequence in the source database may be different from that in the destination database. The replay performance is better. In transaction mode, you cannot set how many concurrent connections are allowed. The number of connections is dynamically adjusted based on the connections in the source database logs to ensure that transaction SQL statements in the same connection of the source database are executed in sequence.
Filter out SQLs	-	The system fuzzily matches SQL statements based on the entered condition.

Parameter	Example Value	Description
Filter out SQLs Without Conditions	-	This option is used to filter out SQL statements of the SELECT, UPDATE, and DELETE types that do not contain conditions.
Maximum Concurrent Connections	8	The number of replay threads configured for the workload replay task.
Acceleration Configuration	100%	The percentage of the replayed SQLs to the SQLs executed on the source database within the same period.

Step 8 Click **Test Connection** in the **Destination Database** area to test the connection to the destination database. After the connection test is successful, click **Next**.

Step 9 On the **Check Task** page, check the workload replay 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 10 On the **Confirm Task** page, specify **Start Time**.

Figure 4-6 Task startup settings

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


★ Send Notifications ?

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

Parameter	Example Value	Description
Start Time	Start upon task creation	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.
Send Notifications	Disable	This parameter is optional. After enabled, select an SMN topic. If the status, latency metric, or data of the workload replay task is abnormal, DRS will send you a notification.
Stop Abnormal Tasks After	14	Any task in the Abnormal state that has run for longer than the period you set here (in days) will automatically stop.

Step 11 Confirm that the configured information is correct and click **Submit** to submit the task.

Step 12 After the workload replay task is submitted, check the task status or [query the replay progress](#) on the **Workload Replay Management** page.

- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, but the task status remains unchanged. When you reconfigure the task, DRS applies for resources for the task again.

----End

Step 2: Query Replay Progress

The replay progress displays the SQL execution status during workload replay, helping you learn about the task status.

Step 1 On the **Workload Replay Management** page, click the target replay task in the **Task Name/ID** column.

Step 2 In the navigation pane, choose **Workload Replay Progress** to check task progress.

- In the **Workload Replay Progress** area, you can view the task status, start time, total number of SQL statements, and number of replayed SQL statements.
- In the **Statistics Chart** area, you can view the total number of SQL statements, number of replayed SQL statements, number of abnormal SQL statements, and number of slow SQL statements in a specified period.
- In the **Abnormal SQLs in Workload Replay** area, you can view the category and number of SQL statements that fail to be replayed.
- In the **Slow SQLs** area, you can view the original time and replay time required for executing a SQL statement.
- In the **SQL Execution Progress** area, you can view the SQL statements that are executing in the destination database during replay.

Step 3 After workload replay is complete, [check the replay reporting](#) to learn about the execution of each SQL statement.

----End

Step 3: Check Replay Reporting

The replay reporting records the execution time curve of each SQL statement replayed in the destination database, number of replayed SQLs, and replay duration.

Step 1 On the **Workload Replay Management** page, click the target replay task in the **Task Name/ID** column.

Step 2 On the **Replay Reporting** page, check the report details about the current task.

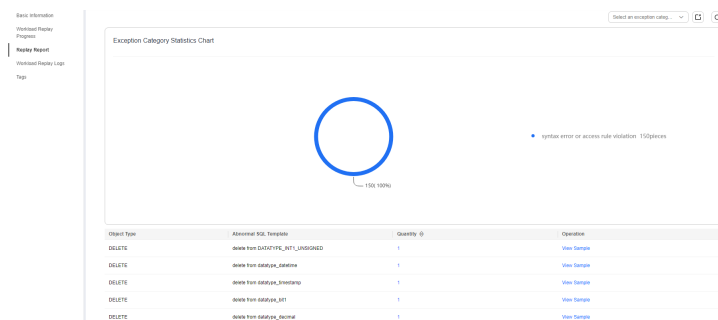
- In the **Statistics Chart** area, you can view the total number of SQL statements, number of replayed SQL statements, number of abnormal SQL statements, and number of slow SQL statements replayed in the destination database in a specified period.

Figure 4-7 Replay statistics chart



- In the **Slow SQLs** area, you can view the number of SQL statements of each type and the replay duration.
- In the **Abnormal SQLs in Workload Replay** area, you can view the category and number of SQL statements that fail to be replayed.

Figure 4-8 SQLs to be replayed



----End

Stopping a Task

After confirming that the workload replay task is no longer used, you can stop the DRS task.

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

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

----End

5 Common Tasks

After completing basic preparations such as accounts, permissions, databases, and networks, you can view common practices to better use DRS.

Table 5-1 Common practices

Scenario		Practice	Description
Creating a Task	Real-Time Migration	From Other Cloud MySQL to RDS for MySQL	This practice describes how to use DRS to migrate data from a MySQL database on another cloud to a Huawei Cloud RDS for MySQL instance through a public network.
		From Other Cloud MySQL to GaussDB(for MySQL)	This practice describes how to use DRS to migrate data from a MySQL database on another cloud to a Huawei Cloud GaussDB(for MySQL) instance through a public network.
		From Other Cloud MongoDB to DDS	This practice describes how to use DRS to migrate data from a MongoDB database on another cloud to a Huawei Cloud DDS instance through a public network.
		From ECS-hosted MySQL to RDS for MySQL	This practice describes how to use DRS to migrate data from a MySQL database built on an ECS to an RDS for MySQL instance in the same VPC of the same region through a VPC.
		From ECS-hosted MySQL to GaussDB(for MySQL)	This practice describes how to use DRS to migrate data from a MySQL database built on an ECS to a GaussDB(for MySQL) instance in the same VPC of the same region through a VPC.

Scenario		Practice	Description
		From ECS-hosted MongoDB to DDS	This practice describes how to use DRS to migrate data from a MongoDB database built on an ECS to a DDS instance in the same VPC of the same region through a VPC.
		From On-Premises MySQL to RDS for MySQL	This practice describes how to use DRS to migrate data from an on-premises MySQL database to a Huawei Cloud RDS for MySQL instance through a public network.
		From On-Premises MongoDB to DDS	This practice describes how to use DRS to migrate data from an on-premises MongoDB database to a Huawei Cloud DDS instance through a public network.
		From RDS for MySQL to DDM	This practice describes how to use DRS to migrate data from a Huawei Cloud RDS for MySQL instance to a DDM instance in different regions through a VPN.
		From MySQL Schema and Logic Table to DDM	This practice describes how to use DRS to migrate data from MySQL shards and tables to a DDM instance through a public network.
Backup Migration		Migrating Microsoft SQL Server Backup Data to RDS for SQL Server	This practice describes how to use DRS to restore local Microsoft SQL Server data backups to an RDS for SQL Server instance. DRS supports full backup migration and full+incremental backup migration.
Real-Time Synchronization		From Other Cloud PostgreSQL to RDS for PostgreSQL	This practice describes how to use DRS to synchronize data from a PostgreSQL database on another cloud to an RDS for PostgreSQL instance through a public network.
		From ECS-hosted PostgreSQL to RDS for PostgreSQL	This practice describes how to use DRS to synchronize data from a PostgreSQL database built on an ECS to an RDS for PostgreSQL instance through a VPC.
		From On-Premises PostgreSQL to RDS for PostgreSQL	This practice describes how to use DRS to synchronize data from an on-premises PostgreSQL database to an RDS for PostgreSQL instance through a public network.

Scenario		Practice	Description
		From On-Premises Oracle to GaussDB	This practice describes how to use DRS to create a full+incremental task to continuously synchronize data from an on-premise Oracle database to a GaussDB instance through a public network.
		From On-Premises Oracle to DDM	This practice describes how to use DRS to create a full+incremental task to continuously synchronize data from an on-premise Oracle database to a DDM instance through a public network.
		From RDS for MySQL to Kafka	This practice describes how to use DRS to create an incremental task to synchronize incremental data from an RDS for MySQL instance to a Kafka instance through a VPC.
	Real-Time DR	Configuring Remote Single-Active DR for an RDS for MySQL Instance Using DRS	This practice describes how to use DRS to synchronize data from an RDS for MySQL instance in the production center to an RDS for MySQL instance in the DR center through a public network to implement data DR between the primary instance and the DR instance across regions.
Querying Task Progress	Real-Time Migration	Querying the Migration Progress	DRS shows the task progress using a progress bar, helping you keep track of the status of a task.
	Real-Time Synchronization	Querying the Synchronization Progress	
	Real-Time DR	Querying the DR Progress	
Comparing Data	Real-Time Migration	Comparing Migration Items	Data comparison allows you to check data consistency between source and destination databases before and after the migration. To minimize the impact on services and shorten the service interruption duration, DRS provides multiple comparison methods.
	Real-Time Synchronization	Comparing Synchronization Items	

Scenario		Practice	Description
	Real-Time DR	Comparing DR Items	
Managing Tasks	Real-Time Migration	Migration Task Life Cycle	During the life cycle of a DRS task, you can edit, pause, reset, resume and stop the task, and modify the flow control mode of the task as required.
	Real-Time Synchronization	Synchronization Task Life Cycle	
	Real-Time DR	DR Task Life Cycle	

More Information

- [Real-Time Migration Overview](#)
- [Backup Migration Overview](#)
- [Real-Time Synchronization Overview](#)
- [Real-Time DR Overview](#)