

Data Replication Service

User Guide

Issue 06
Date 2022-05-30



Copyright © Huawei Technologies Co., Ltd. 2023. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Technologies Co., Ltd.

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

Website: <https://www.huawei.com>

Email: support@huawei.com

Contents

1 Service Overview.....	1
1.1 What Is DRS?.....	1
1.2 Advantages.....	2
1.3 Functions and Features.....	2
1.3.1 Real-Time Migration.....	2
1.3.2 Backup Migration.....	10
1.3.3 Real-Time Synchronization.....	12
1.4 Mapping Data Types.....	23
1.4.1 Oracle->MySQL.....	23
1.4.2 Oracle->PostgreSQL.....	25
1.5 Basic Concepts.....	26
1.6 Security Suggestions.....	28
1.7 Accessing DRS.....	29
1.8 Instructions.....	29
1.9 Compliance Description.....	29
1.10 Related Services.....	29
2 Real-Time Migration.....	31
2.1 To the Cloud.....	31
2.1.1 From MySQL to MySQL.....	31
2.1.2 From Oracle to PostgreSQL.....	48
2.1.3 From MongoDB to DDS.....	52
2.2 Out of the Cloud.....	68
2.2.1 From MySQL to MySQL.....	68
2.2.2 From DDS to MongoDB.....	81
2.3 Task Management.....	93
2.3.1 Creating a Migration Task.....	93
2.3.2 Querying the Migration Progress.....	103
2.3.3 Viewing Migration Logs.....	104
2.3.4 Comparing Migration Items.....	105
2.3.5 Managing Objects.....	108
2.3.5.1 Migrating Accounts.....	108
2.3.5.2 Parameters for Comparison.....	110
2.3.6 Task Life Cycle.....	112

2.3.6.1 Viewing Task Details.....	112
2.3.6.2 Editing Migration Task Information.....	113
2.3.6.3 Modifying Connection Information.....	114
2.3.6.4 Modifying the Flow Control Mode.....	114
2.3.6.5 Editing a Migration Task.....	115
2.3.6.6 Resuming a Migration Task.....	119
2.3.6.7 Resetting a Migration Task.....	120
2.3.6.8 Pausing a Migration Task.....	120
2.3.6.9 Stopping a Migration Task.....	121
2.3.6.10 Deleting a Migration Task.....	122
2.3.6.11 Task Statuses.....	123
2.4 Tag Management.....	124
2.5 Interconnecting with Cloud Eye.....	125
2.5.1 Supported Metrics.....	125
2.5.2 Configuring Alarm Rules.....	130
2.5.3 Viewing Monitoring Metrics.....	130
3 Backup Migration.....	132
3.1 Creating an RDS Backup Migration Task.....	132
3.2 Creating a Backup Using OBS Buckets.....	136
3.3 Task Management.....	142
3.3.1 Viewing Task Details.....	142
3.3.2 Editing Migration Task Information.....	142
3.3.3 Viewing Migration Logs.....	143
3.3.4 Deleting a Migration Task.....	143
3.3.5 Task Statuses.....	144
3.4 Tag Management.....	144
4 Real-Time Synchronization.....	146
4.1 Data Synchronization Topologies.....	146
4.2 To the Cloud.....	148
4.2.1 From MySQL to MySQL.....	148
4.2.2 From PostgreSQL to PostgreSQL.....	163
4.2.3 From Oracle to MySQL.....	180
4.2.4 From Oracle to PostgreSQL.....	197
4.3 Out of the Cloud.....	204
4.3.1 From MySQL to MySQL.....	204
4.3.2 From MySQL to Kafka.....	217
4.4 Task Management.....	225
4.4.1 Creating a Synchronization Task.....	225
4.4.2 Querying the Synchronization Progress.....	233
4.4.3 Viewing Synchronization Logs.....	234
4.4.4 Comparing Synchronization Items.....	235
4.4.5 Managing Objects.....	237

4.4.5.1 Editing Synchronization Objects.....	237
4.4.5.2 Mapping Object Names.....	238
4.4.5.3 Viewing Synchronization Mapping Information.....	239
4.4.5.4 Processing Data.....	240
4.4.6 Task Life Cycle.....	243
4.4.6.1 Viewing Task Details.....	243
4.4.6.2 Modifying Task Information.....	244
4.4.6.3 Modifying Connection Information.....	245
4.4.6.4 Modifying the Flow Control Mode.....	245
4.4.6.5 Editing a Synchronization Task.....	246
4.4.6.6 Resuming a Synchronization Task.....	247
4.4.6.7 Pausing a Synchronization Task.....	248
4.4.6.8 Resetting a Synchronization Task.....	249
4.4.6.9 Stopping a Synchronization Task.....	250
4.4.6.10 Deleting a Synchronization Task.....	251
4.4.6.11 Task Statuses.....	252
4.5 Tag Management.....	253
4.6 Interconnecting with Cloud Eye.....	254
4.6.1 Supported Metrics.....	254
4.6.2 Configuring Alarm Rules.....	259
4.6.3 Viewing Monitoring Metrics.....	259
4.7 Operation Reference in Synchronization Scenarios.....	260
4.7.1 Kafka Message Format.....	260
4.7.2 Kafka Authentication.....	264
4.7.3 Forcibly Stopping Synchronization of PostgreSQL.....	267
4.7.4 Creating Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.....	269
5 FAQs.....	272
5.1 Product Consulting.....	272
5.1.1 What Are Regions and AZs?.....	272
5.1.2 What Is DRS?.....	273
5.1.3 Can DRS Migrate RDS Primary/Standby Instances?.....	274
5.1.4 Does DRS Support Resumable Uploads?.....	274
5.1.5 What Are the Differences Between Real-Time Migration and Real-Time Synchronization?.....	275
5.1.6 How Do I Solve the Table Bloat Issue?.....	276
5.1.7 How Do I Select RDS Read Replicas on the DRS Console?.....	276
5.1.8 How Does DRS Affect the Source and Destination Databases?.....	276
5.1.9 Can DRS Migrates Table Structures Only?.....	277
5.2 Network and Security.....	277
5.2.1 What Security Protection Policies Does DRS Have?.....	277
5.2.2 What Can I Do If the Network Is Disconnected During the Migration?.....	278
5.2.3 How Can I Configure a VPC Security Group to Allow Traffic from an EIP?.....	278

5.2.4 What Can I Do If the Network Connection Between the Replication Instance and Database Is Abnormal?.....	278
5.3 Permissions Management.....	281
5.3.1 Which MySQL Permissions Are Required for DRS?.....	281
5.3.2 How Can I Import Users and Permissions from the Source to the Destination Database?.....	283
5.4 Real-Time Migration.....	284
5.4.1 When Can I Stop a Migration Task?.....	284
5.4.2 How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration?.....	285
5.4.3 What Can I Do If the Invoking Permission Problem Occurs After the MySQL Stored Procedure Is Migrated to the Cloud?.....	286
5.4.4 How Do I Ensure that All Services on the Database Are Stopped?.....	287
5.4.5 What Can I Do When Scheduled Task Failed to Start Leaving Message "can not get agency token".....	288
5.4.6 What Can I Do If MyISAM Tables Are Not Supported by RDS for MySQL?.....	288
5.4.7 What Are the Precautions for Migrating Data from an Earlier Version MySQL to MySQL 8.0?.....	288
5.4.8 What Can I Do When OOM Occurs During the Migration of MongoDB Databases?.....	299
5.4.9 How Do I Disable the Balancer?.....	299
5.4.10 How Do I Export and Import Events and Triggers in Batches?.....	300
5.4.11 How Can I Migrate Databases or Tables Whose Names Contain Uppercase Letters?.....	301
5.4.12 How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?.....	302
5.5 Backup Migration.....	303
5.5.1 What Should I Do If the Last Backup File Is Incorrectly Selected in the Backup Migration Scenario?.....	303
5.5.2 Manual Configuration.....	303
5.6 Real-Time Synchronization.....	311
5.6.1 Can DRS Sync Tables of Different Schemas to the Same Schema?.....	311
5.6.2 Can Online DDL Tools Be Used for Real-time Synchronization?.....	311
5.6.3 Which MySQL DDL Statements Can Be Directly Executed in GaussDB(DWS)?.....	311
5.6.4 Why Do I Use the SCAN IP Address to Connect to an Oracle RAC Cluster?.....	312
5.6.5 How Do I Check Supplemental Logging of the Source Oracle Database?.....	312
5.6.6 Garbled Characters or Synchronization Failure Due to Incompatible Character Sets.....	314
5.7 Data-Level Comparison.....	314
5.7.1 Which of the Following Data Types Are Not Supported By Value Comparison?.....	314
5.7.2 What Impact Does a DRS Comparison Task Have on Databases?.....	315
5.7.3 How Long Does a DRS Comparison Task Take?.....	315
5.8 General Operations.....	316
5.8.1 What Can I Do When Information Overlaps on the DRS Console?.....	316
5.8.2 Is the Destination Instance Set to Read-only or Read/Write?.....	316
5.8.3 How Do I Set Global binlog_format=ROW to Take Effect Immediately?.....	316
5.8.4 How Do I Set binlog_row_image=FULL to Take Effect Immediately?.....	317
5.8.5 How Do I Change the Destination Database Password to Meet the Password Policy?.....	318
5.8.6 How Do I Configure the Shard Key for a MongoDB Sharded Cluster?.....	319
5.8.7 Does Bandwidth Expansion Affect the Running DRS Tasks?.....	321

5.8.8 Why Data in MariaDB and SysDB Cannot Be Migrated?.....	321
5.8.9 Constraints and Operation Suggestions on Many-to-One Scenario.....	321
5.8.10 Where Can I View DRS Operation Logs?.....	323
5.8.11 Can a Completed Task Be Restarted?.....	324
5.8.12 What Are the Differences Between Resetting a Task and Recreating a Task?.....	324
5.8.13 What Do I Do After Changing the Password of the Source or Destination Database?.....	324
5.9 Delay.....	324
5.9.1 Why Does the Delay of DR Tasks Increase?.....	324
5.9.2 Why Is the Delay High In MongoDB Replication Scenarios?.....	326
6 Troubleshooting.....	327
6.1 Solutions to Failed Check Items.....	327
6.1.1 Disk Space.....	327
6.1.1.1 Checking Whether the Destination Database Has Sufficient Storage Space.....	327
6.1.1.2 Checking Whether the Destination Server Has Sufficient Storage Space.....	329
6.1.2 Database Parameters.....	329
6.1.2.1 Checking Whether the Source Database Binlog Is Enabled.....	329
6.1.2.2 Checking Whether the Source Database Binlog Is Row-Based.....	330
6.1.2.3 Checking Whether the Binlog Retention Period Is Set on the Source Database.....	331
6.1.2.4 Checking Whether the Source and Destination Database Character Sets Are Consistent.....	332
6.1.2.5 Checking Whether the Source Database server_id Meets the Incremental Migration Requirements	333
6.1.2.6 Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity.....	334
6.1.2.7 Checking Whether the Source Database Contains Object Names with Non-ASCII Characters.....	335
6.1.2.8 Checking Whether the TIME_ZONE Values of the Source and Destination Databases Are the Same	335
6.1.2.9 Checking Whether the COLLATION_SERVER Values of the Source and Destination Databases Are the Same.....	336
6.1.2.10 Checking Whether the SERVER_UUID Values of the Source and Destination Databases Are the Same.....	336
6.1.2.11 Checking Whether the SERVER_ID Values of the Source and Destination Databases Are Different	337
6.1.2.12 Checking Whether the Source Database Contains Invalid sql_mode Values.....	337
6.1.2.13 Checking Whether the SQL_MODE Values of the Source and Destination Databases Are the Same.....	338
6.1.2.14 Checking Whether the sql_mode Value in the Destination Database Is Not NO_ENGINE_SUBSTITUTION.....	338
6.1.2.15 Checking Whether the innodb_strict_mode Values of the Source and Destination Databases Are the Same.....	339
6.1.2.16 Checking Whether the max_wal_senders Value of the Source Database Is Correctly Configured	340
6.1.2.17 Checking Whether the WAL_LEVEL Value in the Source Database Is Correct.....	340
6.1.2.18 Checking Whether the MAX_REPLICATION_SLOTS Value in the Source Database Is Correct.....	341
6.1.2.19 Checking Whether the Source Database Is on Standby.....	342

6.1.2.20 Checking Whether the log_slave_updates Value of the Source Database Is Correctly Configured	342
6.1.2.21 Checking Whether the BLOCK_SIZE Value of the Source Database Is the Same as That of the Destination Database	343
6.1.2.22 Checking Whether the binlog_row_image Value is FULL	344
6.1.2.23 Checking Whether the Transaction Isolation Levels are Consistent	344
6.1.2.24 Checking Whether the lc_monetary Values of the Source and Destination Databases Are the Same	345
6.1.2.25 Checking Whether the Source Database Contains Trigger Names with Non-ASCII Characters	345
6.1.2.26 Checking Whether the Source Database Collections Contain More Than 10 Indexes	346
6.1.2.27 Checking Whether log_bin_trust_function_creators Is Set to On in Both the Source and Destination Databases	346
6.1.2.28 Checking Whether log_bin_trust_function_creators Is Set to On in the Destination Database	347
6.1.2.29 Checking Whether the max_allowed_packet Value of the Destination Database Is too Small	348
6.1.2.30 Checking Whether the Databases and Tables Exist	348
6.1.2.31 Checking Whether session_replication_role of the Destination Database Is correctly Set	349
6.1.2.32 Checking Whether the MongoDB Instance Type Matches the Migration Mode	349
6.1.2.33 Checking the Physical Standby Database	350
6.1.2.34 Checking Whether the Values of group_concat_max_len Are Consistent	350
6.1.2.35 Checking Whether the Character Sets Are Compatible	351
6.1.2.36 Checking Replication Attribute of Primary Key Columns	351
6.1.2.37 Checking Whether the Source and Destination Database Character Sets Are Consistent	352
6.1.3 Source DB Instance Statuses	353
6.1.3.1 Checking Whether the Source and Destination Databases Are of the Same Type	353
6.1.3.2 Checking Whether the ChangeStream API of the source DB instance is available	353
6.1.4 Destination DB Instance Statuses	354
6.1.4.1 Checking Whether the Destination Database Is Involved in Another Migration Task	354
6.1.4.2 Checking Whether the Destination Database Has a Read Replica	355
6.1.4.3 Checking Whether the Extensions Are Supported	355
6.1.4.4 Checking Whether the Destination DB Instance Is Available	356
6.1.5 Database User Permissions	356
6.1.5.1 Whether the Source Database User Has Sufficient Permissions	356
6.1.5.2 Checking Whether the Destination Database User Has Sufficient Permissions	360
6.1.6 Database Versions	363
6.1.6.1 Checking Whether the Source Database Version Is Supported	363
6.1.6.2 Checking Whether the Destination Database Version Is Supported	364
6.1.6.3 Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version	365
6.1.7 Networks	367
6.1.7.1 Checking Whether the Source Database Is Connected	368
6.1.7.2 Checking Whether the Destination Database Is Connected	370
6.1.7.3 Checking Whether the Destination Database Can Connect to the Source Database	372
6.1.8 Database Objects	372
6.1.8.1 Checking Whether the Source Database Contains a MyISAM Table	373

6.1.8.2 Checking Whether the Source Database Contains the Functions or Stored Procedures that the Source Database User Is Not Authorized to Migrate.....	373
6.1.8.3 Checking Whether the Source Database Tables Use Storage Engines Not Supported by the Destination Database.....	374
6.1.8.4 Checking Whether the Source Database Tables Contain Primary Keys.....	374
6.1.8.5 Checking Whether the Source Database Contains Triggers or Events.....	375
6.1.9 Database Configuration Items.....	375
6.1.9.1 Checking Whether the Source Database Name Is Valid.....	375
6.1.9.2 Checking Whether the Source Database Table Name Is Valid.....	376
6.1.9.3 Checking Whether the Source Database View Name Is Valid.....	376
6.1.9.4 Checking Whether the Shard Key Can Be Obtained from the Source Database.....	377
6.1.10 Conflicts.....	377
6.1.10.1 Checking Whether the Names of the Source and Destination Databases Are the Same.....	377
6.1.10.2 Checking Whether the Destination Database Contains a Non-Empty Collection with the Same Name As the Source Database.....	378
6.1.11 SSL Connections.....	379
6.1.11.1 Checking Whether the SSL Connection Is Correctly Configured.....	379
6.1.11.2 Checking Whether the SSL Connection Is Enabled for the Source Database.....	380
6.1.11.3 Checking Whether the SSL Certificate of the Destination Database Exists.....	380
6.1.12 Object Dependencies.....	381
6.1.12.1 Checking Whether Referenced Tables Are Selected for Migration.....	381
6.2 Failure Cases.....	381
6.2.1 Backup Migration Failed Because Backup Files Cannot Be Found.....	382
6.2.2 Backup Migration Failed Because a Backup Database Cannot Be Found in the Backup Files.....	382
A Change History.....	384

1 Service Overview

This chapter provides general information about Data Replication Service (DRS), including its application scenarios, functions, and constraints.

1.1 What Is DRS?

DRS is a stable, efficient, and easy-to-use cloud service for database migration and synchronization in real time.

It simplifies data migration processes and reduces migration costs.

You can use DRS to quickly transmit data between different DB engines.

Real-Time Migration

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

You can perform a migration task over multiple types of networks, such as public networks, VPCs, VPNs, and direct connections. With these network connections, you can migrate between different cloud platforms, from on-premises databases to cloud databases, or between cloud databases across regions.

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

Backup Migration

It often becomes necessary to hide the 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.

You can use backup migration when you want to migrate on-premises databases to the cloud.

Without connecting to your sources, DRS can help you complete data migration.

Real-Time Synchronization

Real-time synchronization refers to the real-time flow of key service data from sources to destinations through a synchronization instance while consistency of data can be ensured.

It is different from migration. Migration means moving your overall database from one platform to another. Synchronization refers to the continuous flow of data between different services.

You can use real-time synchronization in many scenarios such as real-time analysis, report system, and data warehouse environment.

Real-time synchronization is mainly used for synchronizing tables and data. It can meet various requirements, such as many-to-one, one-to-many synchronization, dynamic addition and deletion of tables, and synchronization between tables with different names.

1.2 Advantages

Easy to Use

DRS simplifies migration procedures and does not require too much technical knowledge. Traditional migration requires professional technical personnel and migration procedures are complicated.

Fast Setup

DRS sets up a migration task within minutes. Traditional migration takes several days, weeks, or even months to set up.

Low Costs

DRS saves traditional database administrator (DBA) labor costs and hardware costs.

Secure

DRS allows you to query the migration progress, check migration logs, and compare migration items, so you can easily complete migration tasks.

1.3 Functions and Features

1.3.1 Real-Time Migration

In real-time migration, you only need to configure the source database, destination database, and migration objects. DRS will help you compare and

analyze data so you can determine when to migrate with minimal service disruption.

Supported Database Types

DRS supports migration between different DB engines. The following table lists the supported data sources. Self-built databases include on-premises databases and ECS databases.

Table 1-1 Database types

Migration Direction	Data Flow	Source DB	Destination DB	Destination DB Type
To the cloud	MySQL->MySQL	<ul style="list-style-type: none">On-premises databasesECS databasesDatabases on other cloudsRDS MySQL DB instances	RDS MySQL DB instances	<ul style="list-style-type: none">Single DB instancePrimary/Standby DB instance
To the cloud	MongoDB->DDS	<ul style="list-style-type: none">On-premises databasesECS databasesDatabases on other cloudsDDS DB instances	DDS DB instances	<ul style="list-style-type: none">ClusterReplica setSingle node
To the cloud	Oracle->PostgreSQL	<ul style="list-style-type: none">On-premises databasesECS databases	RDS PostgreSQL DB instances NOTE The destination DB instance can only be an RDS PostgreSQL Enhanced Edition instance.	<ul style="list-style-type: none">Single DB instancePrimary/Standby DB instance

Migration Direction	Data Flow	Source DB	Destination DB	Destination DB Type
From the cloud	MySQL->MySQL	RDS MySQL DB instances	<ul style="list-style-type: none"> On-premises databases ECS databases Databases on other clouds 	<ul style="list-style-type: none"> Single DB instance Primary/Standby DB instance
From the cloud	DDS->MongoDB	DDS DB instances	<ul style="list-style-type: none"> On-premises databases ECS databases Databases on other clouds 	<ul style="list-style-type: none"> Cluster Replica set Single node

Table 1-2 Database versions

Migration Direction	Data Flow	Source DB Version	Destination DB Version
To the cloud	MySQL->MySQL	<ul style="list-style-type: none"> MySQL 5.5.x MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x 	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x
To the cloud	MongoDB->DDS	<ul style="list-style-type: none"> MongoDB 3.2.x MongoDB 3.4.x MongoDB 4.0.x 	<ul style="list-style-type: none"> DDS 3.2.x DDS 3.4.x DDS 4.0.x DDS 4.2.x <p>NOTE If the destination database version is 4.2, the source database version cannot be later than 4.0.</p>

Migration Direction	Data Flow	Source DB Version	Destination DB Version
To the cloud	Oracle->PostgreSQL	<ul style="list-style-type: none"> Oracle 10g Oracle 11g Oracle 12c Oracle 18c 	PostgreSQL Enhanced Edition
From the cloud	MySQL->MySQL	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x 	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x
From the cloud	DDS->MongoDB	<ul style="list-style-type: none"> DDS 3.2.x DDS 3.4.x DDS 4.0.x 	<ul style="list-style-type: none"> MongoDB 3.2.x MongoDB 3.4.x MongoDB 4.0.x

Supported Migration Types

DRS supports two migration types: full migration and full+incremental migration.

This full 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 that can be migrated include tables, views, stored procedures, and triggers.

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.

Table 1-3 Migration types

Migration Direction	Data Flow	Full Migration	Full+Incremental Migration
To the cloud	MySQL->MySQL	Supported	Supported

Migration Direction	Data Flow	Full Migration	Full+Incremental Migration
To the cloud	MongoDB->DDS	<ul style="list-style-type: none"> • Replica set -> Single node • Replica set -> Replica set • Replica set -> Cluster • Single node -> Single node • Single node -> Replica set • Single node -> Cluster • Cluster -> Cluster 	<ul style="list-style-type: none"> • Replica set -> Single node • Replica set -> Replica set • Replica set -> Cluster • Single node -> Single node • Single node -> Replica set • Single node -> Cluster <p>NOTE</p> <ul style="list-style-type: none"> • If you need to perform an incremental migration for a single-node instance, the source database must be the DDS single-node instance. • The source database cannot be a GaussDB(for Mongo) instance.
To the cloud	Oracle->PostgreSQL	Supported	Supported
From the cloud	MySQL->MySQL	Supported	Supported
From the cloud	DDS->MongoDB	Supported	Supported NOTE If the source database is on a cluster instance, incremental migration is not supported.

Supported Network Types

DRS supports data migration through a Virtual Private Cloud (VPC), Virtual Private Network (VPN), Direct Connect, or public network. [Table 1-4](#) lists the application scenarios of each network type and required preparations, and [Table 1-5](#) lists the supported network types of each migration scenario.

Table 1-4 Network types

Network Type	Application Scenario	Preparations
VPC	Migrations between cloud databases in the same region	<ul style="list-style-type: none">• The source and destination databases must be in the same region.• The source and destination databases can be in either the same VPC or in different VPCs.• If source and destination databases are in the same VPC, they can communicate with each other by default. Therefore, you do not need to configure a security group.• If the source and destination databases are not in the same VPC, the CIDR blocks of the source and destination databases cannot be duplicated or overlapped, and the source and destination databases are connected through a VPC peering connection. DRS automatically establishes a route through a single IP address when you test the network connectivity.
VPN	Migrations from on-premises databases to cloud databases or between cloud databases across regions	Establish a VPN connection between your local data center and the VPC that hosts the destination database. Before migration, ensure that the VPN network is accessible.
Direct Connect	Migrations from on-premises databases to cloud databases or between cloud databases across regions	Use a dedicated network connection to connect your data center to VPCs.

Network Type	Application Scenario	Preparations
Public network	Migrations from on-premises or other cloud databases to destination databases	<p>To ensure network connectivity between the source and destination databases, perform the following operations:</p> <ol style="list-style-type: none"> 1. Enable public accessibility. Enable public accessibility for the source database based on your service requirements. 2. Configure security group rules. <ul style="list-style-type: none"> • Add the EIPs of the replication instance to the whitelist of the source database for inbound traffic. • If destination databases and the replication instance are in the same VPC, they can communicate with each other by default. You do not need to configure a security group. <p>NOTE</p> <ul style="list-style-type: none"> • The IP address on the Configure Source and Destination Databases page is the EIP of the replication instance. • If SSL is not enabled, migrating confidential data is not recommended.

Table 1-5 Supported network types

Migration Direction	Data Flow	VPC	Public Network	VPN or Direct Connect
To the cloud	MySQL->MySQL	Supported	Supported	Supported
To the cloud	MongoDB->DDS	Supported	Supported	Supported
To the cloud	Oracle->PostgreSQL	Not supported	Supported	Supported
From the cloud	MySQL->MySQL	Supported	Supported	Supported

Migration Direction	Data Flow	VPC	Public Network	VPN or Direct Connect
From the cloud	DDS->MongoDB	Supported	Supported	Supported

Migration Objects

DRS allows you to migrate objects at different levels. The following table lists the supported migration objects.

Table 1-6 Supported migration objects

Migration Direction	Data Flow	Full Migration	Table-Level Migration	Database-Level Migration
To the cloud	MySQL->MySQL	Supported	Supported	Supported
To the cloud	MongoDB->DDS	Supported	Supported	Supported
To the cloud	Oracle->PostgreSQL	Not supported	Supported	Not supported
From the cloud	MySQL->MySQL	Supported	Supported	Supported
From the cloud	DDS->MongoDB	Supported	Supported	Supported

Advanced Features

DRS supports multiple features to ensure successful real-time migration.

Table 1-7 Advanced features

Feature	Description
Flow control	Allows you to limit the overall migration speed to make the impact of migration on bandwidth and database I/O controllable. Flow control mode takes effect only during a full migration.
Account migration	Allows you to migrate accounts, permissions, and passwords.
Parameter comparison	Checks the consistency of common parameters and performance parameters between source and destination databases to ensure that the migrated service is running properly.

1.3.2 Backup Migration

DRS supports backup migrations of various database types.

Supported Database Types

Table 1-8 Database types

Data Flow	Backup File Source	Destination DB Type
Microsoft SQL Server -> RDS for SQL Server	<ul style="list-style-type: none">On-premises Microsoft SQL Server backup filesRDS for SQL Server full backup filesMicrosoft SQL Server backup files on other clouds	RDS for SQL Server DB instances

Migration Methods

Table 1-9 Migration methods

Data Flow	Full Migration	Incremental Migration
Microsoft SQL Server -> RDS for SQL Server	Supported	Supported

Supported Database Versions

Table 1-10 Database versions

Data Flow	Backup File Version	Destination DB Version
Microsoft SQL Server -> RDS for SQL Server	On-premises and other cloud's Microsoft SQL Server backup file versions: <ul style="list-style-type: none"> • Microsoft SQL Server 2000 • Microsoft SQL Server 2005 • Microsoft SQL Server 2008 • Microsoft SQL Server 2012 • Microsoft SQL Server 2014 • Microsoft SQL Server 2016 • Microsoft SQL Server 2017 • Microsoft SQL Server 2019 	<ul style="list-style-type: none"> • Microsoft SQL Server 2008 • Microsoft SQL Server 2012 • Microsoft SQL Server 2014 • Microsoft SQL Server 2016 • Microsoft SQL Server 2017 • Microsoft SQL Server 2019
	RDS for SQL Server full backup file versions: <ul style="list-style-type: none"> • Microsoft SQL Server 2008 • Microsoft SQL Server 2012 • Microsoft SQL Server 2014 • Microsoft SQL Server 2016 • Microsoft SQL Server 2017 • Microsoft SQL Server 2019 	<ul style="list-style-type: none"> • Microsoft SQL Server 2008 • Microsoft SQL Server 2012 • Microsoft SQL Server 2014 • Microsoft SQL Server 2016 • Microsoft SQL Server 2017 • Microsoft SQL Server 2019

Backup Migration Scenarios

Table 1-11 Migration scenarios

Scenario	Description
OBS bucket	If you copy the database backup files to an Object Storage Service (OBS) bucket, ensure that the OBS bucket is located in the same region as the destination instance.
RDS full backup	If you select an RDS full backup as the backup file source, ensure that the RDS instance has a full backup.

1.3.3 Real-Time Synchronization

Real-time synchronization refers to the real-time flow of key service data from sources to destinations while consistency of data can be ensured. It is different from migration. Migration means moving your overall database from one platform to another. Synchronization refers to the continuous flow of data between different services.

Supported Database Types

DRS supports real-time synchronization between databases of various types, and many-to-one synchronization.

Table 1-12 Database types

Sync roniza tion Dire ction	Data Flow	Source DB	Destination DB	Destina tion DB Type
To the clou d	MySQL->MySQL	<ul style="list-style-type: none">On-premises databasesECS databasesDatabases on other cloudsRDS MySQL DB instances	RDS MySQL DB instances	<ul style="list-style-type: none">Single DB instancePrimary/Standby DB instance

Sync hron izati on Dire ction	Data Flow	Source DB	Destination DB	Destina tion DB Type
To the clou d	PostgreSQL- >PostgreSQL	<ul style="list-style-type: none"> • On-premises databases • ECS databases • Databases on other clouds • RDS PostgreSQL DB instances 	RDS PostgreSQL DB instances	<ul style="list-style-type: none"> • Singl e DB insta nce • Prim ary/ Stan dby DB insta nce
To the clou d	Oracle->PostgreSQL	<ul style="list-style-type: none"> • On-premises databases • ECS databases 	RDS PostgreSQL DB instances	<ul style="list-style-type: none"> • Singl e DB insta nce • Prim ary/ Stan dby DB insta nce
To the clou d	Oracle->MySQL	<ul style="list-style-type: none"> • On-premises databases • ECS databases 	RDS MySQL DB instances	<ul style="list-style-type: none"> • Singl e DB insta nce • Prim ary/ Stan dby DB insta nce

Sync hron izati on Dire ction	Data Flow	Source DB	Destination DB	Destina tion DB Type
From the clou d	MySQL->MySQL	RDS MySQL DB instances	<ul style="list-style-type: none">• On-premises databases• ECS databases• Databases on other clouds• RDS MySQL DB instances	-
From the clou d	MySQL->Kafka	RDS MySQL DB instances	Kafka	<ul style="list-style-type: none">• Clust er• Singl e node

Synchronization Methods

DRS supports three synchronization modes: full synchronization, incremental synchronization, and full+incremental synchronization.

Full synchronization: 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 synchronization: Through log parsing, DRS replicates incremental data to keep sources and destinations in sync.

Full+Incremental synchronization: DRS 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.

Table 1-13 Synchronization methods

Sync hron izati on Dire ction	Data Flow	Incremental	Full	Full +Incre mental	One- way/ Two- way Sync
To the clou d	MySQL->MySQL	Supported	Not support ed	Support ed	One-way sync
To the clou d	PostgreSQL- >PostgreSQL	Supported	Support ed	Support ed	One-way sync
To the clou d	Oracle->PostgreSQL	Not supported	Support ed	Support ed	One-way sync
To the clou d	Oracle->MySQL	Not supported	Support ed	Support ed	One-way sync
From the clou d	MySQL->MySQL	Supported	Not support ed	Support ed	One-way sync
From the clou d	MySQL->Kafka	Supported	Not support ed	Not support ed	One-way sync

Database Versions

 **NOTE**

Data cannot be synchronized from a newer version database to an older version database.

Table 1-14 Database versions

Sync roniza tion Dire ctio n	Data Flow	Source Database Version	Destination DB Version
To the clou d	MySQL->MySQL	<ul style="list-style-type: none"> ● MySQL 5.5.x ● MySQL 5.6.x ● MySQL 5.7.x ● MySQL 8.0.x 	<ul style="list-style-type: none"> ● MySQL 5.6.x ● MySQL 5.7.x ● MySQL 8.0.x
To the clou d	PostgreSQL->PostgreSQL	<ul style="list-style-type: none"> ● PostgreSQL 9.4.x ● PostgreSQL 9.5.x ● PostgreSQL 9.6.x ● PostgreSQL 10.x ● PostgreSQL 11.x ● PostgreSQL 12.x ● PostgreSQL 13.x <p>NOTE If the source DB type is RDS PostgreSQL Enhanced Edition, the destination DB type must be the same.</p>	<ul style="list-style-type: none"> ● PostgreSQL 9.5.x ● PostgreSQL 9.6.x ● PostgreSQL 10.x ● PostgreSQL 11.x ● PostgreSQL 12.x ● PostgreSQL 13.x
To the clou d	Oracle->PostgreSQL	<ul style="list-style-type: none"> ● Oracle 10g ● Oracle 11g ● Oracle 12c ● Oracle 18C ● Oracle 19C 	<ul style="list-style-type: none"> ● PostgreSQL 9.5.x ● PostgreSQL 9.6.x ● PostgreSQL 10.x ● PostgreSQL 11.x ● PostgreSQL 12.x
To the clou d	Oracle-> MySQL	<ul style="list-style-type: none"> ● Oracle 10g ● Oracle 11g ● Oracle 12c ● Oracle 18C ● Oracle 19C 	<ul style="list-style-type: none"> ● MySQL 5.6.x ● MySQL 5.7.x ● MySQL 8.0.x

Sync hron izati on Dire ctio n	Data Flow	Source Database Version	Destination DB Version
From the clou d	MySQL->MySQL	<ul style="list-style-type: none"> • MySQL 5.6.x • MySQL 5.7.x • MySQL 8.0.x 	<ul style="list-style-type: none"> • MySQL 5.6.x • MySQL 5.7.x • MySQL 8.0.x
From the clou d	MySQL->Kafka	<ul style="list-style-type: none"> • MySQL 5.6.x • MySQL 5.7.x 	Kafka 0.11 or later
Self- built -> Self- built	Oracle -> GaussDB primary/standby	<ul style="list-style-type: none"> • Oracle 10g • Oracle 11g • Oracle 12c • Oracle 18c • Oracle 19c 	GaussDB 1.0.0 or later
Self- built -> Self- built	Oracle -> GaussDB distributed	<ul style="list-style-type: none"> • Oracle 10g • Oracle 11g • Oracle 12c • Oracle 18c • Oracle 19c 	GaussDB 1.0.0 or later

Network Types

DRS supports real-time synchronization through a Virtual Private Cloud (VPC), Virtual Private Network (VPN), Direct Connect, or public network. [Table 1-15](#) lists the application scenarios of each network type and required preparations.

Table 1-15 Network types

Network Type	Application Scenario	Preparations
VPC	Synchronization between cloud databases in the same region	<ul style="list-style-type: none">• The source and destination databases must be in the same region.• The source and destination databases can be in either the same VPC or in different VPCs.• If source and destination databases are in the same VPC, they can communicate with each other by default. Therefore, you do not need to configure a security group.• If the source and destination databases are not in the same VPC, the CIDR blocks of the source and destination databases cannot be duplicated or overlapped, and the source and destination databases are connected through a VPC peering connection. DRS automatically establishes a route through a single IP address when you test the network connectivity.
VPN	Synchronization from on-premises databases to cloud databases or between cloud databases across regions	Establish a VPN connection between your local data center and the VPC that hosts the destination database. Before synchronization, ensure that the VPN network is accessible.
Direct Connect	Synchronization from on-premises databases to cloud databases or between cloud databases across regions	Use a dedicated network connection to connect your data center to VPCs.

Network Type	Application Scenario	Preparations
Public network	Synchronization from on-premises or external cloud databases to the destination databases.	<p>To ensure network connectivity between the source and destination databases, perform the following operations:</p> <ol style="list-style-type: none"> 1. Enable public accessibility. Enable public accessibility for the source database based on your service requirements. 2. Configure security group rules. <ul style="list-style-type: none"> • Add the EIPs of the synchronization instance to the whitelist of the source database for inbound traffic. • If destination databases and the synchronization instance are in the same VPC, they can communicate with each other by default. Therefore, you do not need to configure a security group. <p>NOTE</p> <ul style="list-style-type: none"> • The IP address on the Configure Source and Destination Databases page is the EIP of the synchronization instance. • If SSL is not enabled, synchronizing confidential data is not recommended.

Table 1-16 Supported network types

Synchronization Direction	Data Flow	VPC	Public Network	VPN or Direct Connect
To the cloud	MySQL->MySQL	Supported	Supported	Supported

Synchr onizatio n Dire ctio n	Data Flow	VPC	Public Network	VPN or Direct Connect
To the clou d	PostgreSQL->PostgreSQL	Supported	Supported	Supported
To the clou d	Oracle->MySQL	Supported	Supported	Supported
To the clou d	Oracle->PostgreSQL	Supported	Supported	Supported
Fro m the clou d	MySQL->MySQL	Supported	Supported	Supported
Fro m the clou d	MySQL->Kafka	Supported	Supported	Supported

Supported Synchronization Objects

DRS allows you to synchronize different objects. The following table lists the supported objects.

Table 1-17 Supported synchronization objects

Synchr onizatio n Dire ctio n	Data Flow	Table- level	Database- level	Importing an Object File
To the clou d	MySQL->MySQL	Supported	Supported	Supported
To the clou d	PostgreSQL->PostgreSQL	Supported	Supported	Supported
To the clou d	Oracle->MySQL	Supported	Not supported	Supported
To the clou d	Oracle->PostgreSQL	Supported	Not supported	Supported
Fro m the clou d	MySQL->MySQL	Supported	Supported	Not supported
Fro m the clou d	MySQL->Kafka	Supported	Supported	Supported

Advanced Features

DRS supports multiple features to ensure successful data synchronization.

Table 1-18 Advanced features

Feature	Description
Synchronization level	<p>DRS supports database- and table-level synchronization.</p> <ul style="list-style-type: none">• Database-level synchronization refers to a type of synchronization method using database as a unit. You do not need to select tables to be synchronized. New tables in the database are automatically added to the synchronization task.• Table-level synchronization uses table as a unit, indicating that you need to add new tables to the synchronization task manually.
Mapping object names	<p>Allows the names of synchronization objects (including databases, schemas, tables, and columns) in the source database to be different from those in the destination database. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one.</p> <p>The following objects can be mapped: databases, schemas and tables.</p>
Dynamically adding or deleting synchronization objects	<p>During data synchronization, you can add or delete synchronization objects as required.</p>

Feature	Description
Conflict policy	<p>DRS uses primary key or unique key conflict policies to ensure that tables with primary key or unique constraints in the source database can be synchronized to the destination database as expected.</p> <p>The following conflict policies are supported:</p> <ul style="list-style-type: none"> • Ignore The system will skip the conflicting data and continue the subsequent synchronization process. • Overwrite Conflicting data will be overwritten. • Report error The synchronization task will be stopped and fail. <p>Ignore and overwrite: Synchronization stability is prioritized, so tasks will not be interrupted as data conflicts occur.</p> <p>Report error: Data quality is prioritized. Any data conflicts are not allowed, so once a conflict occurs, the synchronization task fails and an error is reported. You need to manually find the cause of the fault. If the task is in the failed state for a long time, the storage space may be used up and the task cannot be restored.</p>
Structure synchronization	DRS does not provide data structure synchronization as an independent function during real-time synchronization. Instead, it directly synchronizes data and structures to the destination database.

1.4 Mapping Data Types

DRS allows you to migrate or synchronize between sources and destinations that use different DB engines through mappings between different data types.

This section provides mappings between different data types for your reference.

1.4.1 Oracle->MySQL

Table 1-19 Data type mapping

Data Type (Oracle)	Condition	Data Type (MySQL)	Whether to Support Mapping
CHAR	length<=255	CHAR	Yes
CHAR	length>255	VARCHAR	Yes

Data Type (Oracle)	Condition	Data Type (MySQL)	Whether to Support Mapping
VARCHAR	Length (row size) ≤ 65536	VARCHAR	Yes
VARCHAR	Length (row size) > 65536	TEXT	Yes
VARCHAR2	-	VARCHAR2	Yes
NCHAR	length≤255	NCHAR	Yes
NCHAR	length>255	NVARCHAR	Yes
NVARCHAR2	-	NVARCHAR	Yes
NUMBER	precision=0 scale = 0	DECIMAL(65,30)	Yes
NUMBER	precision!=0 scale!=0	DECIMAL(precision, scale)	Yes
FLOAT	-	FLOAT	Yes
BINARY_FLOAT	-	FLOAT	Yes
BINARY_DOUBLE	-	DOUBLE	Yes
DATE	-	DATETIME	Yes
TIMESTAMP	-	DATETIME	Yes
TIMESTAMP WITH TIME ZONE	6 digit precision	TIMESTAMP	Yes
TIMESTAMP WITH LOCAL TIME ZONE	6 digit precision	TIMESTAMP	Yes
INTERVAL	6 digit precision	VARCHAR(30)	Yes
BLOB	-	LOB	Yes
CLOB	-	LONGTEXT	Yes
NCLOB	-	LONGTEXT	Yes
LONG	-	LONGTEXT	Yes
RAW	-	VARBINARY	Yes
LONG RAW	-	LOB	Yes
ROWID	-	VARCHAR(18)	Yes
UROWID	-	-	No
XMLTYPE	-	-	No

Data Type (Oracle)	Condition	Data Type (MySQL)	Whether to Support Mapping
BFILE	-	-	No
SDO_GEOMETRY	-	-	No

1.4.2 Oracle->PostgreSQL

Oracle -> PostgreSQL Community Edition

Table 1-20 Data type mapping

Data Type (Oracle)	Data Type (PostgreSQL Community Edition)	Whether to Support Mapping
CHAR	CHAR	Yes
VARCHAR	VARCHAR	Yes
VARCHAR2	VARCHAR	Yes
NCHAR	NCHAR	Yes
NVARCHAR2	VARCHAR	Yes
NUMBER	NUMBER	Yes
BINARY_FLOAT	FLOAT	Yes
BINARY_DOUBLE	DOUBLE	Yes
FLOAT	FLOAT	Yes
DATE	TIMESTAMP	Yes
TIMESTAMP	TIMESTAMP	Yes
TIMESTAMP WITH TIME ZONE	TIMESTAMPTZ	Yes
TIMESTAMP WITH LOCAL TIME ZONE	TIMESTAMPTZ	Yes
INTERVAL	INTERVAL	Yes
BLOB	BYTEA	Yes
CLOB	CLOB	Yes
NCLOB	TEXT	Yes
LONG	TEXT	Yes
LONG_RAW	BYTEA	Yes

Data Type (Oracle)	Data Type (PostgreSQL Community Edition)	Whether to Support Mapping
RAW (non-primary key and non-unique key column)	BYTEA	Yes
RAW (primary key and unique key column)	VARCHAR	Yes
ROWID	CHAR	Yes
UROWID	-	No
XMLTYPE	-	No
BFILE	-	No
SDO_GEOMETRY	-	No

1.5 Basic Concepts

VPC

VPC-based migration refers to a real-time migration that the source and destination databases are in the same VPC or two VPCs that can communicate with each other. No additional network services are required.

VPN

VPN-based migration refers to a real-time migration where the source and destination databases are in the same VPN network. The VPN establishes a secure, encrypted communication tunnel that complies with industry standards between your data centers and the cloud platform. Through this tunnel, DRS seamlessly migrates data from the data center to the cloud.

Direct Connect

Direct Connect enables you to establish a dedicated network connection from your data center to the cloud platform. With Direct Connect, you can use a dedicated network connection to connect your data center to VPCs to enjoy a high-performance, low-latency, and secure network.

Replication Instance

A replication instance refers to an instance that performs the migration task. It exists in the whole lifecycle of a migration task. DRS uses the replication instance to connect to the source database, read source data, and replicate the data to the destination database.

Migration Log

A migration log refers to the log generated during database migration. Migration logs are classified into the following levels: warning, error, and info.

Task Check

Before starting a migration task, you need to check whether the source and destination databases have met all migration requirements. If any check item fails, rectify the fault and check the task again. Only when all check items are successful the task can start.

To the Cloud

DRS requires that either the source or destination database is on the current cloud. **To the cloud** means that the destination database must be on the current cloud.

Out of the Cloud

DRS requires that either the source or destination database is on the current cloud. **Out of the cloud** means that the source database must be on the current cloud.

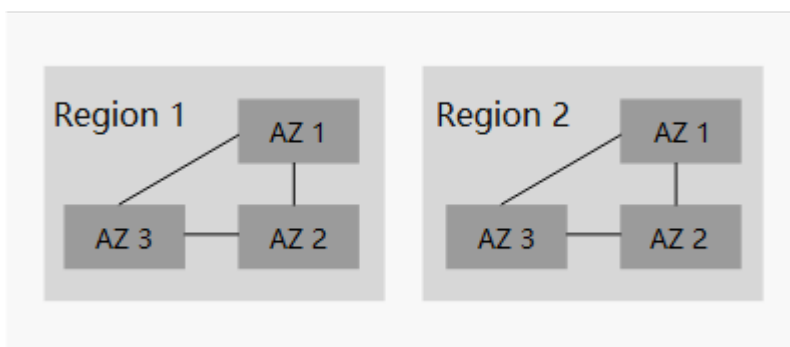
Region and AZ

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center. Each region is completely independent, improving fault tolerance and stability. After a resource is created, its region cannot be changed.
- An AZ is a physical location using independent power supplies and networks. Faults in an AZ do not affect other AZs. A region can contain multiple AZs, which are physically isolated but interconnected through internal networks. This ensures the independence of AZs and provides low-cost and low-latency network connections.

Figure 1-1 shows the relationship between regions and AZs.

Figure 1-1 Region and AZ



Account Entrustment

DRS will entrust your account to the administrator to implement some functions. For example, if you enable scheduled startup tasks, DRS will automatically entrust your account to DRS administrator **op_svc_rds** during the task creation to implement automated management on the scheduled tasks.

Account entrustment can be implemented in the same region only.

Temporary Accounts

To ensure that your database can be successfully migrated to RDS MySQL DB instances, DRS automatically creates temporary accounts **drsFull** and **drsIncremental** in the destination database during full migration and incremental migration, respectively. After the migration task is complete, DRS automatically deletes the temporary account.

NOTICE

Attempting to delete, rename, or change the passwords or permissions for temporary accounts will cause task errors.

High Availability

If the primary host of a replication instance fails, it automatically fails over to the standby host, preventing service interruption and improving the success rate of migration.

If a replication instance fails, the system will automatically restart the instance and retry the task. In this case, the task status changes to **Fault rectification**. If the replication instance is still faulty after being restarted, the system automatically creates an instance. After the instance is created, the system retries the task again. The high availability management applies to the following tasks:

- Full migration
- Incremental migration

1.6 Security Suggestions

You can encrypt data and database connections and increase database account security.

Data Encryption

Before saving sensitive information and private data in a database, encrypt data to reduce the risk of information leakage.

Connection Encryption

Use the Secure Sockets Layer (SSL) to encrypt connections between applications and DB instances to enhance data transmission security.

Database Account Security

To ensure data security, improve the security of database accounts.


1.7 Accessing DRS

Prerequisites

To begin using DRS, register an account on the official website. After the registration is successful, you can access all cloud services, including DRS, RDS, and DDS.

If you have registered an account, you can log in to the management console and access your DRS.

Procedure

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Log in to the management console, click **Data Replication Service** under **Databases** to go to the DRS console.

----End

1.8 Instructions

To improve your experience with DRS, add the Security Administrator permission using IAM in case some functions become unavailable, such as scheduled task startup, automatic ending of full-migration tasks, and automatic retry of failed tasks.

1.9 Compliance Description

You understand and agree that your use of this service complies with laws and regulations, including but not limited to legal compliance requirements on data content, data transfer, and cross-border data transfer. DRS only provides a standard service upon your request and is not responsible for the legal compliance of your use. If you use the services illegally or engage in illegal actions using the services, you shall bear all consequences arising therefrom.

1.10 Related Services

RDS

DRS can migrate data from your databases to the RDS databases in the cloud. For more information about RDS, see Relational Database Service User Guide.

Supported network types during migration to RDS:

- VPC
- VPN
- Direct Connect
- Public network

2 Real-Time Migration

2.1 To the Cloud

2.1.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 2-1 Supported databases

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

Supported Migration Objects

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

Table 2-2 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">• Object level: table level, database level, or instance level (full migration).• Supported migration objects:<ul style="list-style-type: none">– Databases, tables, users, views, indexes, constraints, functions, stored procedures, triggers, and events– The system database and event statuses cannot be migrated.– Tables with storage engine different to MyISAM and InnoDB tables cannot be migrated.– Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys– Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. <p>NOTE The objects that can be migrated have the following constraints:</p> <ul style="list-style-type: none">• The source database name, table name, and view name cannot contain non-ASCII characters or special characters '<>`\'\"'• The source database name cannot start with ib_logfile and cannot be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1.

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-3](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

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

Table 2-3 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT The user must have the SELECT permission for mysql.user.	The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT The user must have the SELECT permission for mysql.user.
Destination database user	The user must have the following minimum permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database.	

Suggestions

CAUTION

- When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To maintain data consistency before and after the migration, do not write data to the source and destination databases in the full migration mode. In the full +incremental migration mode, you can continue the migration while data is still being written to the source database.
-
- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
 - Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration. If you have to migrate data during peak hours, you can select **Yes** for **Flow Control** to adjust the migration speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full migration, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be migrated without a primary key may be locked for 3s.
 - The data being migrated may be locked by other transactions for a long period of time, resulting in read timeout.

- Due to the inherent characteristics of MySQL, in some scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
- If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- Data-level comparison
To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

Precautions

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

Table 2-4 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none">● Source database parameter requirements:<ul style="list-style-type: none">- The binlog of the source database must be enabled and use the row-based format during incremental migration.- If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the migration may fail. If the source database is an on-premises MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide.- During an incremental migration, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296.- Enable skip-name-resolve for the source database to reduce the possibility of connection timeout.- Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.- The log_slave_updates parameter of the source database must be enabled. Otherwise, the migration fails.- The binlog_row_image parameter value of the source database must be FULL. Otherwise, the migration fails.● Source database object requirements:<ul style="list-style-type: none">- If the source database is an on-premises database and has Percona Server for MySQL 5.6.x or Percona Server for MySQL 5.7.x installed, the memory manager must use Jemalloc to prevent Out of Memory errors caused by frequent queries on system tables.- The source database does not support the mysql binlog dump command.- The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.- Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.

Type	Restrictions
	<ul style="list-style-type: none"> - Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index. - Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database. • Destination database parameter requirements: <ul style="list-style-type: none"> - The destination DB instance is running properly. - The destination DB instance must have sufficient storage space. - You are advised to use the row-based binlog in the destination database. Otherwise, an error may occur during an incremental migration. - The destination database isolation level must be set to at least read committed. - During migration, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB. - Enable GTID of the destination database. - The character sets of the source and destination databases must be the same. Otherwise, the migration fails. - The log_bin_trust_function_creators parameter value of the destination database must be set to on. Otherwise, the migration fails. - The value of server_uuid of the destination database must be the same as that of the source database. Otherwise, the incremental migration fails. - The collation_server values of the destination database and source database must be the same. Otherwise, the migration may fail. - The value of time_zone of the destination database must be the same as that of the source database. Otherwise, the migration may fail. - The sql_mode values of the destination database and source database must be the same. Otherwise, the migration may fail. - The innodb_strict_mode values of the destination database and source database must be the same. Otherwise, the migration may fail.

Type	Restrictions
	<ul style="list-style-type: none">- The lower_case_table_names values of the source and destination databases must be the same. Otherwise, the migration fails.- If the MyISAM tables are included in the migration objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the migration fails.● Destination database object requirements:<ul style="list-style-type: none">- The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).● Other notes:<ul style="list-style-type: none">- When creating multiple migration tasks in the many-to-one scenario, ensure that the read and write settings of the destination database are consistent in these tasks.- The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.- The destination database cannot be restored to a point in time when a full migration was being performed.- If the source and destination sides are RDS MySQL instances, transparent data encryption (TDE) is not supported, and tables with the encryption function cannot be created.- If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.- The destination database of a migration task can be set to Read-only or Read/Write. Read-only: During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration. Read/Write: During the migration, the destination instance can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks.

Type	Restrictions
Full migration	<ul style="list-style-type: none">• During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.• During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated.• During migration, do not write the statement-based binlog into the source database.• During migration, do not clear the binlog in the source database.• During migration, do not create a database named ib_logfile in the source database.
Incremental migration	<ul style="list-style-type: none">• During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated.• During migration, do not write the statement-based binlog into the source database.• During migration, do not clear the binlog in the source database.• During migration, do not create a database named ib_logfile on the source side.• During an incremental migration of table-level objects, renaming tables is not supported.• During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database.• During an incremental migration, resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down.• DDL statements are supported in the incremental migration phase.• If the source and destination databases are of the same major version and the entire instance is migrated, DCL statements can be migrated in incremental mode, but users cannot be changed by updating the mysql.user table. For details about DCL statements, see the MySQL official document.

Type	Restrictions
Stopping a task	<ul style="list-style-type: none">• Stop a task normally.<ul style="list-style-type: none">– The selected events and triggers are migrated while the migration task proceeds to the final stage. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the migration status reported by the migration log.• Forcibly stop a task.<ul style="list-style-type: none">– If you forcibly stop a task, DRS resources will be released and triggers and events will not be migrated. You need to manually migrate triggers and events. If you want DRS to migrate triggers and events, restore the DRS task first. After the task status becomes normal, stop the task.

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time migration, see [Real-Time Migration](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

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

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

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

- Task information description

Table 2-5 Task information

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

- Replication instance information

Table 2-6 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination DB is on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	Select VPC Network. Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network . <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region.- Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN and Direct Connect are suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions.
Destination DB Instance	The RDS DB instance you created.
Replication Instance Subnet	The subnet where the replication instance resides. You can also click View Subnet to go to the network console to view the subnet where the instance resides. By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Destination Database Access	<ul style="list-style-type: none">- Read-only During migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.- Read/Write During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks. Set the destination database to Read/Write only when you need to modify other data in the database during the migration. The task cannot be modified after being created.
Migration Type	<ul style="list-style-type: none">- Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects and data in non-system databases are migrated to the destination database at one time. The objects include tables, views, and stored procedures. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.- Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.

- Enterprise Projects and Tags

Table 2-7 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">- Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.- After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 On the **Configure Source and Destination Databases** page, wait until the replication instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

 **NOTE**

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

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

Table 2-8 Self-build on ECS - source database information

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

Parameter	Description
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is not enabled, your data may be at risk.

 **NOTE**

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

- Scenario 2: RDS DB instance - source database configuration

Table 2-9 RDS DB instance - source database information

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

- Destination database configuration

Table 2-10 Destination database settings

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

Parameter	Description
Migrate Definer to User	<ul style="list-style-type: none"><li data-bbox="683 297 1442 566">– 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?<li data-bbox="683 577 1442 712">– No The Definers of all source database objects will not be changed. You need to migrate all accounts and permissions of the source database in the next step.


 **NOTE**

The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Task** page, select the accounts and objects to be migrated, and click **Next**.

Table 2-11 Migration types and objects

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The migration speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none">• Yes If you need to migrate accounts, see Migrating Accounts.• No During migration, accounts, permissions, and passwords are not migrated.

Parameter	Description
Filter DROP DATABASE	<p>To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations.</p> <ul style="list-style-type: none">• If you select Yes, any database deletion operations performed on the source database are not migrated during data migration.• If you select No, related operations are migrated to the destination database during data migration.
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none">• All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified.• Tables: The selected table-level objects will be migrated.• Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none">• If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected migration object cannot contain spaces.• To quickly select the desired database objects, you can use the search function.

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

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

For details about how to handle check items that fail to pass the pre-check, see [Solutions to Failed Check Items](#).

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

 **NOTE**

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

Step 6 Compare source and destination parameters.

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

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

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

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

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

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

 **NOTE**

You can also manually enter parameter values.

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

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

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

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


- iii. Click **Next**.

Step 7 On the displayed page, specify **Start Time** and confirm that the configured information is correct and click **Submit** to submit the task.

Table 2-12 Task startup settings

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

Step 8 After the task is submitted, view and manage it on the **Online Migration Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

2.1.2 From Oracle to PostgreSQL

Supported Source and Destination Databases

Table 2-13 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises Oracle databases• Oracle databases on an ECS	<ul style="list-style-type: none">• RDS for PostgreSQL

Suggestions

CAUTION

- When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To maintain data consistency before and after the migration, do not write data to the source and destination databases in the full migration mode. In the full +incremental migration mode, you can continue the migration while data is still being written to the source database.
-

- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
- Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full migration, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be migrated without a primary key may be locked for 3s.
 - The data being migrated may be locked by other transactions for a long period of time, resulting in read timeout.
 - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- Data-level comparison
To obtain accurate comparison results, compare data at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, inconsistent comparison results may be generated, reducing the reliability and validity of the results.

Precautions

Before creating a migration task, read the following notes:

Table 2-14 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none">• The source database must have the following minimum permissions: CREATE SESSION, SELECT ANY TRANSACTION, SELECT ANY TABLE, SELECT ANY DICTIONARY and SELECT ANY SEQUENCE.• The destination database must have the following minimum permissions: The permissions to create databases and tables in the destination database.

Type	Restrictions
Migration objects	<ul style="list-style-type: none"> • Tables, indexes, constraints, sequences, and data can be migrated. Other database objects, such as stored procedures, cannot be migrated. • Full synchronization does not support the following column types: bfile, xml, sdo_geometry, urowid, and self-defined types. • For the TIMESTAMP WITH TIME ZONE data type, the data cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database. • Tables whose default values contain expressions of functions cannot be migrated. • The default values of the source database can be to_date and sys_guid functions. To use a function as the default value, ensure that the destination database has the same function. For example, if the source database contains the sys_guid function, the uuid_generate_v1 function is required in the destination PostgreSQL database. If the sys_guid function fails to be migrated, run the following statement in the destination database to install the plug-in and retry the task: create extension if not exists "uuid-osspl"; • Function indexes cannot be migrated. Example: create index idx_t on t(substr(dt, 1, 8));
Source database	<ul style="list-style-type: none"> • PostgreSQL and Oracle database structures are different. The former has a schema. Therefore the statement for creating views cannot contain db.Table. Otherwise, the view migration fails. Example: Change the statement 1 to statement 2. Statement 1: create view v1 as select id from db1.t1; Statement 2: create view v1 as select id from t1; • The maximum precision supported by timestamp and interval day to second is 6. • The data type cannot be bfile, xmltype, sdo_geometry, or user-defined. • The source database cannot contain tables with same names but different letter cases. • The names of databases and tables cannot contain non-ASCII characters or special characters .><` ,?!" • If the destination database is PostgreSQL Enhanced Edition, the trigger name in the source database must be unique. • An empty source database cannot be migrated. • Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, and WE8MSWIN1252.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• The destination instance must be the RDS PostgreSQL Enhanced Edition DB instance.• The destination database cannot contain objects whose names are lowercase version of the source database object names.
Precautions	<ul style="list-style-type: none">• Objects that have dependencies must be migrated at the same time to avoid migration failure. Common dependencies: tables referenced by primary or foreign keys• You can migrate one database (owner) at a time in a single migration task. To migrate multiple databases, you need to create multiple tasks.• After objects such as tables are migrated to the destination database, their names are converted into lowercase letters. For example, the name of object ABC is changed to abc after being migrated to the destination database.• The time zone settings of the source and destination database must be the same.• If there are special characters in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database.• The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.• If the data types are incompatible, the migration may fail.• During the migration, writing data to the source and destination databases is not allowed. Otherwise, data inconsistency may occur.• During the migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• During task startup or full migration, you are not advised to perform DDL operations on the source database, such as deleting the database or indexes. Otherwise, the migration may fail.• During a full migration, DRS writes large amount of data to the destination PostgreSQL database. As a result, the number of PostgreSQL WAL logs increases sharply, and the PostgreSQL disk space may be used up. You can disable the PostgreSQL log backup function before the full migration to reduce the number of WAL logs. After the migration is complete, enable the function. <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</p>

2.1.3 From MongoDB to DDS

Supported Source and Destination Databases

Table 2-15 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">On-premises Mongo database (versions 3.2, 3.4, and 4.0)Self-built MongoDB on ECS (versions 3.2, 3.4, and 4.0)MongoDB 3.2, 3.4, and 4.0 on other clouds (Tencent Cloud MongoDB 3.2 is not supported.)DDS DB instances (versions 3.4 and 4.0)	<ul style="list-style-type: none">DDS DB instance (versions 3.4, 4.0, and 4.2) <p>NOTE The destination database version must be the same as or later than the source database version.</p>

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-16](#). DRS will automatically check the objects you selected before the migration.

Table 2-16 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">• Object level: table level, database level, or instance level (full migration).• Supported migration objects:<ul style="list-style-type: none">- Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. Common dependencies: collections referenced by views, and views referenced by views- Replica set: Only collections (including validator and capped collections), indexes, and views can be migrated.- Cluster: Only collections (including validator and capped collections), shard keys, indexes, and views can be migrated.- Single node: Only collections (including validator and capped collections), indexes, and views can be migrated.- Only user data and source database account information can be migrated. The system databases (for example, local, admin, and config) and system collection cannot be migrated. If service data is stored in the system database, run the renameCollection command to move the service data to the user database.- The statement for creating a view cannot contain a regular expression.- Collections that contain the _id field without indexes are not supported.- The first parameter of BinData() cannot be 2.- If ranged sharding is used, maxKey cannot be used as the primary key. <p>NOTE The objects that can be migrated have the following constraints:</p> <ul style="list-style-type: none">• The source database name cannot contain \., \$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-17](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 2-17 Database account permission

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

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

 **NOTE**

For example, the source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database.

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

Suggestions

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the migration to help you detect and resolve problems in advance.
- In the migration, ensure that no data is written to the destination database to ensure data consistency before and after the migration.
- Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full migration, and two to four CPUs are occupied.
 - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- Data-Level Comparison
To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

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

Table 2-18 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none"> ● Source database parameter requirements: <ul style="list-style-type: none"> - During an incremental migration, Olog of the source database must be enabled. - If the storage space is sufficient, store the source database Olog for as long as possible. The recommended retention period is three days. - If the incremental source data cluster is migrated, the source database balancer must be disabled. ● Source database object requirements: <ul style="list-style-type: none"> - The source cannot be a GaussDB(for Mongo) instance. - The source database name cannot contain /\.,\$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$). - When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique. ● Destination database parameter requirements: <ul style="list-style-type: none"> - The destination DB instance is running properly. - The destination DB instance must have sufficient storage space. - The version of the destination database instance must be the same as or later than that of the source database. - DRS supports full migration between cluster instances. If the source cluster instance is not sharded, ensure that the size of the primary shard on the destination database is greater than that of the source database. ● Other notes: <ul style="list-style-type: none"> - Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. Common dependencies: collections referenced by views, and views referenced by views - Replica set: The MongoDB replica set instance must be available and have primary nodes. - Source database from a single node instance on other clouds cannot be migrated. - If you need to perform an incremental migration for a single node instance, the source database must be a DDS single-node instance on the current cloud. - If the source database is not on a cluster instance, the following operations and commands are supported during incremental migration: Creating and deleting databases

Type	Restrictions
	<p>Adding, deleting, and updating documents</p> <p>Creating and deleting collections</p> <p>Creating and deleting indexes</p> <p>Creating and deleting views</p> <p>The <code>convertToCapped</code>, <code>collMod</code>, and <code>renameCollection</code> commands are supported.</p> <ul style="list-style-type: none"> - During a full plus incremental migration between clusters, the objects to be migrated cannot be deleted. Otherwise, the migration task fails. - If you select Cluster (MongoDB 4.0+) for Source DB Instance Type, DRS will use the MongoDB change streams feature during the migration. Note the following before you use change streams: Data subscription using change streams consumes a certain amount of CPU and memory resources of the source database. Evaluate the resources of the source database in advance. If the load on the source database is heavy, the processing speed of change streams cannot keep up with the oplog generation speed. As a result, DRS synchronization delay occurs. Change streams support only the following DDLs: drop database, drop collection and rename The DBPointer and DBRef data types are not supported. In the incremental migration phase, the migration speed can reach up to 10,000 rows in a single table per second. Currently, only whitelisted users can use Change Streams. - If you select Olog for Source Database Type, DRS will create multiple subtasks based on the number of source shards. If Task Rate Limit is selected, the configured rate limit value is synchronized to each subtask. - If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental migration, data consistency cannot be ensured when source and destination databases are in different time zone. - The value of block_compressor is determined by stats().wiredTiger.creationString.block_compressor of the collection in the source database. If the destination database contains corresponding empty collections, the compression parameters will not be migrated. If the compression parameters in the source database are not supported by the destination database, configure the compression parameters based on net.compression.compressors of the destination database.

Type	Restrictions
	<p>If the destination database version is DDS 4.2, DRS does not migrate compression parameters because the destination database does not support compression parameter settings.</p> <ul style="list-style-type: none">- If the accounts and roles to be migrated conflict with those in the destination database, DRS will skip the conflict data and continue the migration.- If the MongoDB service of the source database is deployed with other services on the same server, set the value of the cacheSizeGB parameter to the half of the minimum idle cache for the WiredTiger engine of the source database.- If the source is a replica set instance, enter information about all primary and secondary nodes to reduce the impact of a primary/secondary switchover on the migration task. If you enter information about multiple primary and secondary nodes, ensure that all nodes belong to the same replica set instance.- If the source is a cluster instance, enter information about multiple mongos nodes to reduce the impact of single-node failure on the migration task. In addition, ensure that all mongos nodes belong to the same cluster instance. For an incremental migration of a cluster instance, you are advised to enter information about all primary and secondary nodes of the shard node and ensure that all node information belongs to the same shard to reduce the impact of a primary/secondary switchover on the migration task. Ensure that all shard nodes belong to the same cluster.- During the creation of a migration task, the destination database can be set to read-only or read/write if the DDS kernel version supports this setting. After the task is created, this setting cannot be changed. Read-only: During the migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration. Read/Write: During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks.- In some migration scenarios, to prevent the drop database operation from deleting the existing collections in the destination database, the drop database operation will not be synchronized to the destination database.

Type	Restrictions
	<p>If the source database version is earlier than MongoDB 3.6, running the drop database command will delete the collections only from the source database. The collections in destination database will not be deleted.</p> <p>If the source database version is MongoDB 3.6 or later, the drop database operation is represented by the drop database and drop collection operations in oplog. Running the drop database command will delete the collections from both the source and destination databases.</p> <ul style="list-style-type: none">- To accelerate the migration, delete unnecessary indexes from the source database and retain only necessary indexes before the migration. You are advised not to create indexes for the source database during the migration. If indexes must be created, create them in the background.- To prevent loopback, do not start tasks that migrate the same database to and out of the cloud at the same time.
Full migration	<ul style="list-style-type: none">● During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.● During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.● During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated.● During the migration, data rollback caused by primary/standby switchover of the source database is not supported.● During the migration, documents larger than 16 MB in the source database cannot be inserted or updated.

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

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time migration, see [Real-Time Migration](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section uses MongoDB sharded clusters as an example to describe how to configure a task for migrating MongoDB databases to DDS over a public network.

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

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

- Task information description

Table 2-19 Task information

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

- Replication instance information

Table 2-20 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database must be a database on the current cloud.
Source DB Engine	Select MongoDB .
Destination DB Engine	Select DDS .
Network Type	Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network . <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region.- Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN and Direct Connect are suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions.
Destination DB Instance	Select the DB instance you have created.

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

Parameter	Description
Obtain Incremental Data	<p>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.</p> <ul style="list-style-type: none">- oplog: For MongoDB 3.2 or later, DRS directly connects to each shard of the source DB instance to extract data. If you select this method, you must disable the balancer of the source database. For details, see How Do I Disable the Balancer? When testing the connectivity between the source and the DRS instance, you need to enter the connection information of each shard in the source database on the task configuration page.- changeStream: This method is recommended. For MongoDB 4.0 and later, DRS connects to mongos nodes of the source database to extract data. If you select this method, you must enable the WiredTiger storage engine of the source database. <p>NOTE Only whitelisted users can use changeStream. To use this function, submit a service ticket. In the upper right corner of the management console, choose Service Tickets > Create Service Ticket to submit a service ticket.</p>
Destination Database Access	<p>This option does not appear if the DDS kernel version does not support this setting. This setting cannot be changed after a migration task is created.</p> <ul style="list-style-type: none">- Read-only During migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.- Read/Write During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks. Set the destination database to Read/Write only when you need to modify other data in the database during the migration. <p>The task cannot be modified after being created.</p>

Parameter	Description
Source Shard Quantity	If Source DB Instance Type is set to Cluster and Obtain Incremental Data is set to oplog , you need to enter the number of source database shards. The number of source shards ranges from 2 to 32. Specify this parameter based on the actual number of shards in the source DB.

- Enterprise Projects and Tags

Table 2-21 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">– Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.– After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 On the **Configure Source and Destination Databases** page, wait until the replication instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

- Source database information

Table 2-22 Source database information

Parameter	Description
mongos Address	<p>IP address or domain name of the source database in the IP address/Domain name:Port format. The port of the source database. Range: 1 - 65535</p> <p>You can enter a maximum of three groups of IP addresses or domain names of the source database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080. Ensure that the entered IP addresses or domain names belong to the same sharded cluster.</p> <p>NOTE If multiple IP addresses or domain names are entered, the test connection is successful as long as one IP address or domain name is accessible. Therefore, you must ensure that the IP address or domain name is correct.</p>
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is admin .
mongos Username	The username for accessing the source database.
mongos Password	The password for the database username.
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is not enabled, your data may be at risk.
Sharded Database	Enter the information about the sharded databases in the source database.

NOTE

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

- Destination database configuration

Table 2-23 Destination database settings

Parameter	Description
DB Instance Name	The DB instance you selected when creating the migration task. This parameter cannot be changed.

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


 **NOTE**

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

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

Table 2-24 Migrate Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The migration speed is not limited and the outbound bandwidth of the source database is maximally used, which causes read consumption on the source database accordingly. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- If you select Cluster for Source DB Instance Type and oplog for Obtain Incremental Data in Step 2, the flow control settings will be synchronized to each subtask.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

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

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

- If any check fails, review the cause and rectify the fault. After the fault is rectified, click **Check Again**.
- If the check is complete and the check success rate is 100%, click **Next**.

 **NOTE**


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

- Step 6** On the displayed page, specify **Start Time** and confirm that the configured information is correct and click **Submit** to submit the task.

Table 2-25 Task startup settings

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

- Step 7** After the task is submitted, view and manage it on the **Online Migration Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

2.2 Out of the Cloud

2.2.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 2-26 Supported databases

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

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-27](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

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

Table 2-27 Database account permission

Type	Full Migration	Full+Incremental Migration
Source database user	The user must have the following minimum permissions: SELECT, SHOW VIEW, and EVENT During user migration, the user must have the SELECT permission for mysql.user.	The user must have the following minimum permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT During user migration, the user must have the SELECT permission for mysql.user.
Destination database user	The user must have the following minimum permissions: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required. To migrate data, you must have the SELECT, INSERT, UPDATE, and DELETE permissions for the MySQL database.	

Suggestions

⚠ CAUTION

- When a task is being started or in the full migration phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To maintain data consistency before and after the migration, do not write data to the source and destination databases in the full migration mode. In the full+incremental migration mode, you can continue the migration while data is still being written to the source database.
-
- The success of migration depends on environment and manual operations. You can run a migration test before you start the full-scale migration to help you detect and resolve problems in advance.
 - Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration. If you have to migrate data during peak hours, you can select **Yes** for **Flow Control** to adjust the migration speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full migration, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be migrated without a primary key may be locked for 3s.
 - The data being migrated may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in some scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - Data-level comparison

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

Precautions

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

Table 2-28 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none">● Source database parameter requirements:<ul style="list-style-type: none">- The binlog of the source database must be enabled and use the row-based format during incremental migration.- If the storage space is sufficient, store the source database binlog files as long as possible. The recommended retention period is three days. If you set the period to 0, the migration may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide.- During an incremental migration, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296.- Enable skip-name-resolve for the source database to reduce the possibility of connection timeout.- Enable GTID for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.- The log_slave_updates parameter of the source database must be enabled. Otherwise, the migration fails.- The binlog_row_image parameter value of the source database must be FULL. Otherwise, the migration fails.● Source database object requirements:<ul style="list-style-type: none">- The source database does not support the mysql binlog dump command.- The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.- Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects.- Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.- Due to the MySQL constraints, if the one-time event triggering time of the source database is earlier than the migration start time, the event will not be migrated to the destination database.● Destination database parameter requirements:

Type	Restrictions
	<ul style="list-style-type: none">- The destination DB instance is running properly.- The destination DB instance must have sufficient storage space.- You are advised to use the row-based binlog in the destination database. Otherwise, an error may occur during an incremental migration.- The destination database isolation level must be set to at least read committed.- During migration, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.- Enable GTID of the destination database.- The character sets of the source and destination databases must be the same. Otherwise, the migration fails.- The log_bin_trust_function_creators parameter value of the destination database must be set to on. Otherwise, the migration fails.- The value of server_uuid of the destination database must be the same as that of the source database. Otherwise, the incremental migration fails.- The collation_server values of the destination database and source database must be the same. Otherwise, the migration may fail.- The value of time_zone of the destination database must be the same as that of the source database. Otherwise, the migration may fail.- The sql_mode values of the destination database and source database must be the same. Otherwise, the migration may fail.- The innodb_strict_mode values of the destination database and source database must be the same. Otherwise, the migration may fail.- The lower_case_table_names values of the source and destination databases must be the same. Otherwise, the migration fails. <ul style="list-style-type: none">● Destination database object requirements:<ul style="list-style-type: none">- The destination DB instance cannot contain databases with the same name as the source databases (except the MySQL system database).● Other notes:<ul style="list-style-type: none">- The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to

Type	Restrictions
	<p>retry the task several times, or data inconsistency may occur.</p> <ul style="list-style-type: none"> - The destination database cannot be restored to a point in time when a full migration was being performed. - If the source DB instance is an RDS for MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized. - If the destination MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.
Full migration	<ul style="list-style-type: none"> ● During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. ● During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. ● During migration, do not write the statement-based binlog into the source database. ● During migration, do not clear the binlog in the source database. ● During migration, do not create a database named ib_logfile in the source database.

Type	Restrictions
Incremental migration	<ul style="list-style-type: none">• During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.• During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated.• During migration, do not write the statement-based binlog into the source database.• During migration, do not clear the binlog in the source database.• During migration, do not create a database named ib_logfile on the source side.• During an incremental migration of table-level objects, renaming tables is not supported.• During an incremental migration, do not perform the point-in-time recovery (PITR) operation on the source database.• During an incremental migration, resumable upload is supported. However, data may be repeatedly inserted into a non-transactional table that does not have a primary key when the server system breaks down.• DDL statements are supported in the incremental migration phase.
Stopping a task	<ul style="list-style-type: none">• Stop a task normally.<ul style="list-style-type: none">- The selected events and triggers are migrated while the migration task proceeds to the final stage. Before a task is completed, ensure that the source and destination databases are connected and pay attention to the migration status reported by the migration log.• Forcibly stop a task.<ul style="list-style-type: none">- If you forcibly stop a task, DRS resources will be released and triggers and events will not be migrated. You need to manually migrate triggers and events. If you want DRS to migrate triggers and events, restore the DRS task first. After the task status becomes normal, stop the task.

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time migration, see [Real-Time Migration](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

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

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

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

- Task information description

Table 2-29 Task information

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

- Replication instance information

Table 2-30 Replication instance settings

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

Parameter	Description
Replication Instance Subnet	<p>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.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Migration Type	<ul style="list-style-type: none"> - Full: This migration type is suitable for scenarios where service interruption is acceptable. All objects and data in non-system databases are migrated to the destination database at one time. The objects include tables, views, and stored procedures. <p>NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.</p> <ul style="list-style-type: none"> - Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. <p>NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.</p>

- Enterprise Projects and Tags

Table 2-31 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 On the **Configure Source and Destination Databases** page, wait until the replication instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Table 2-32 Source database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during migration task creation. This parameter cannot be changed.
Database Username	Enter the username of the source database.
Database Password	The password for the database username. If the task is in the Starting , Full migration , Incremental migration , or Incremental migration failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.

 NOTE

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

Table 2-33 Destination database settings

Parameter	Description
VPC	A dedicated virtual network in which the destination database is located. It isolates networks for different services.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	Enter the IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535

Parameter	Description
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full migration, Incremental migration, or Incremental migration failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none">The maximum size of a single certificate file that can be uploaded is 500 KB.If SSL is not enabled, your data may be at risk.
Migrate Definer to User	<ul style="list-style-type: none">Yes The Definers of all source database objects will be migrated to the user. Other users do not have permissions for database objects unless these users are authorized. For details on authorization, see How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration?No The Definers of all source database objects will not be changed. You need to migrate all accounts and permissions of the source database in the next step.


 **NOTE**

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

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

Table 2-34 Migration types and objects

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">● Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.● No The migration speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Filter DROP DATABASE	<p>During an incremental migration, executing DDL operations on the source database may affect the data migration performance to some extent. To reduce data migration risks, DRS allows you to filter out DDL operations.</p> <p>The database deletion operation can be filtered out by default.</p> <ul style="list-style-type: none">● If you select Yes, any database deletion operations performed on the source database are not synchronized during data migration.● If you select No, related operations are synchronized to the destination database during data migration. <p>NOTE</p> <p>Currently, only the full plus incremental migrations from RDS MySQL to MySQL are supported.</p>

Parameter	Description
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements.</p> <ul style="list-style-type: none">• Yes If you need to migrate accounts, see Migrating Accounts.• No During the migration, accounts, permissions, and passwords are not migrated.
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none">• All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified.• Tables: The selected table-level objects will be migrated.• Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none">• If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected migration object cannot contain spaces.• To quickly select the desired database objects, you can use the search function.

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

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

For details about how to handle check items that fail to pass the pre-check, see [Solutions to Failed Check Items](#).

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

 **NOTE**


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

Step 6 On the displayed page, specify **Start Time** and confirm that the configured information is correct and click **Submit** to submit the task.

Table 2-35 Task startup settings

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

Step 7 After the task is submitted, view and manage it on the **Online Migration Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

2.2.2 From DDS to MongoDB

Supported Source and Destination Databases

Table 2-36 Supported Databases

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

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see [Table 2-37](#). DRS will automatically check the objects you selected before the migration.

Table 2-37 Migration objects

Type	Precautions
Migration objects	<ul style="list-style-type: none">• Object level: table level, database level, or instance level (full migration).• Supported migration objects:<ul style="list-style-type: none">- Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. Common dependencies: collections referenced by views, and views referenced by views- Replica set: Only collections (including validator and capped collections), indexes, and views can be migrated.- Cluster: Only collections (including validator and capped collections), shard keys, indexes, and views can be migrated.- Single node: Only collections (including validator and capped collections), indexes, and views can be migrated.- System databases (such as local, admin, and config) cannot be migrated. You need to manually create usernames and roles in the destination database.- The statement for creating a view cannot contain a regular expression.- Collections that contain the <code>_id</code> field without indexes are not supported.- The first parameter of <code>BinData()</code> cannot be <code>2</code>.

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see [Table 2-38](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 2-38 Database account permission

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

 NOTE

For example, the source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database.

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

Suggestions

- The success of database migration depends on environment and manual operations. To ensure a smooth migration, perform a migration trial before you start the migration to help you detect and resolve problems in advance.
- In the migration, ensure that no data is written to the destination database to ensure data consistency before and after the migration.
- Start your migration task during off-peak hours. A less active database is easier to migrate successfully. If the data is fairly static, there is less likely to be any severe performance impacts during the migration.
 - If network bandwidth is not limited, the query rate of the source database increases by about 20 MB/s during full migration, and two to four CPUs are occupied.
 - If DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- Data-Level Comparison
To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

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

Table 2-39 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none">● Source database parameter requirements:<ul style="list-style-type: none">- During an incremental migration, Olog of the source database must be enabled.- If the storage space is sufficient, store the source database Olog for as long as possible. The recommended retention period is three days.- If the incremental source data cluster is migrated, the source database balancer must be disabled.● Source database object requirements:<ul style="list-style-type: none">- The source database name cannot contain \\. \$ or spaces. The collection name and view name cannot start with system. or contain the dollar sign (\$).- When multiple source databases are migrated to the same destination database, the name of the database to be migrated must be unique.● Destination database parameter requirements:<ul style="list-style-type: none">- The destination DB instance is running properly.- The destination DB instance must have sufficient storage space.- The version of the destination database instance must be the same as or later than that of the source database.● Other notes:<ul style="list-style-type: none">- Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. Common dependencies: collections referenced by views, and views referenced by views- If a Time-to-Live (TTL) index already exists in the collection of the source database or is created during an incremental migration, data consistency cannot be ensured when source and destination databases are in different time zone.- The value of block_compressor is determined by stats().wiredTiger.creationString.block_compressor of the collection in the source database. If the destination database contains corresponding empty collections, the compression parameters will not be migrated. If the compression parameters in the source database are not supported by the destination database, configure the compression parameters based on net.compression.compressors of the destination database. If the storage engine of the destination database is not WiredTiger, DRS does not migrate compression parameters.- If the destination database is on a replica set instance out of the cloud, enter information about all primary and

Type	Restrictions
	<p>secondary nodes to reduce the impact of a primary/secondary switchover on the migration task. If you enter information about primary and secondary nodes, ensure that all nodes belong to the same replica set instance.</p> <ul style="list-style-type: none"> - If the destination database is on a cluster instance out of the cloud, you are advised to enter information about multiple mongos nodes to reduce the impact of the single-node failure on the migration task. In addition, multiple mongos nodes support load balancing. In addition, ensure that all mongos nodes belong to the same cluster instance. - To accelerate the migration, delete unnecessary indexes from the source database and retain only necessary indexes before the migration. You are advised not to create indexes for the source database during the migration. If indexes must be created, create them in the background. - To prevent loopback, do not start tasks that migrate the same database to and out of the cloud at the same time.
Full migration	<ul style="list-style-type: none"> ● During task startup and full migration, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. ● During migration, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During migration, do not modify the destination database (including but not limited to DDL and DML operations) that is being migrated. ● During the migration, data rollback caused by primary/standby switchover of the source database is not supported. ● During the migration, documents larger than 16 MB in the source database cannot be inserted or updated.

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

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time migration, see [Real-Time Migration](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

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

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

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

- Task information description

Table 2-40 Task information

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

- Replication instance information

Table 2-41 Replication instance settings

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

Parameter	Description
Replication Instance Subnet	<p>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.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Migration Type	<ul style="list-style-type: none">- Full: This migration type is suitable for scenarios where service interruption is permitted. It migrates all objects and data in non-system databases to the destination database at one time. The objects include collections and indexes. NOTE If you are performing a full migration, do not perform operations on the source database. Otherwise, data generated in the source database during the migration will not be synchronized to the destination database.- Full+Incremental: This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full migration will be continuously synchronized to the destination database, and the source remains accessible.

- Enterprise Projects and Tags

Table 2-42 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">- Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.- After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

- Step 3** On the **Configure Source and Destination Databases** page, wait until the replication instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

Table 2-43 Source database settings

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

 **NOTE**

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

Table 2-44 Destination database settings

Parameter	Description
IP Address or Domain Name	<p>IP address or domain name of the destination database in the IP address/Domain name:Port format. The port of the destination database. Range: 1 - 65535</p> <p>You can enter a maximum of three groups of IP addresses or domain names of the destination database. Separate multiple values with commas (,). For example: 192.168.0.1:8080,192.168.0.2:8080. Ensure that the entered IP addresses or domain names belong to the same instance.</p> <p>NOTE</p> <p>If multiple IP addresses or domain names are entered, the test connection is successful as long as one IP address or domain name is accessible. Therefore, you must ensure that the IP address or domain name is correct.</p>
Authentication Database	The name of the authentication database. For example: The default authentication database of DDS instance is admin .

Parameter	Description
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none">The maximum size of a single certificate file that can be uploaded is 500 KB.If SSL is not enabled, your data may be at risk.


 **NOTE**

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

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

Table 2-45 Migrate Object

Parameter	Description
Flow Control	You can choose whether to control the flow. <ul style="list-style-type: none">Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.No The migration speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. NOTE <ul style="list-style-type: none">Flow control mode takes effect only during a full migration.You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

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

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

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

For details about how to handle check items that fail to pass the pre-check, see [Solutions to Failed Check Items](#).

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

 **NOTE**


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

Step 6 On the displayed page, specify **Start Time** and confirm that the configured information is correct and click **Submit** to submit the task.

Table 2-46 Task startup settings

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

Step 7 After the task is submitted, view and manage it on the **Online Migration Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

2.3 Task Management

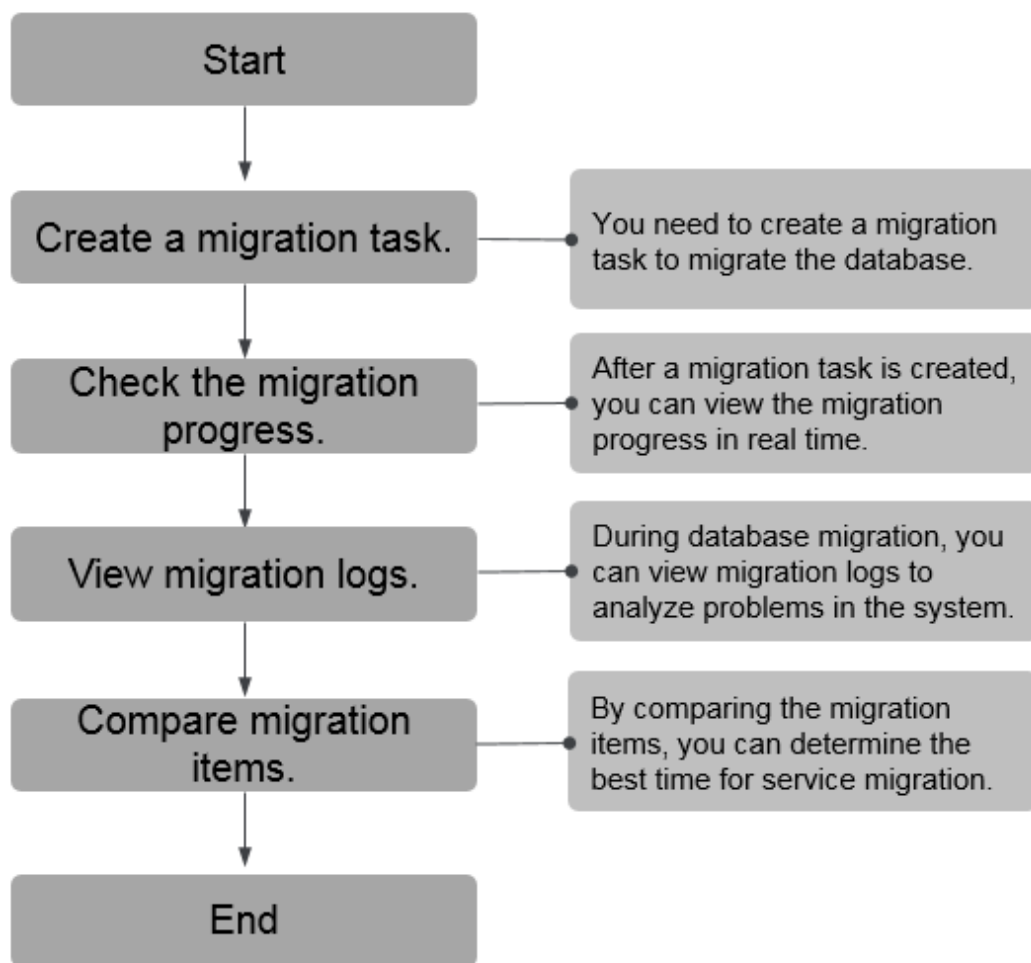
2.3.1 Creating a Migration Task

Process

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

A complete migration involves the following procedures.

Figure 2-1 Migration process



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

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

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

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

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time migration, see [Real-Time Migration](#).

Procedure

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

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

- Task information description

Table 2-47 Task information

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

- Replication instance information

Table 2-48 Replication instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination DB is on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	Select VPC Network. Available options: VPC , VPN or Direct Connect , and Public network . By default, the value is Public network . <ul style="list-style-type: none">- VPC is suitable for migrations between cloud databases of the same account in the same region.- Public network is suitable for migrations from on-premises or external cloud databases to the destination databases bound with an EIP.- VPN and Direct Connect are suitable for migrations from on-premises databases to cloud databases or between cloud databases across regions.

Parameter	Description
Destination DB Instance	The RDS DB instance you created.
Replication Instance Subnet	<p>The subnet where the replication instance resides. You can also click View Subnet to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides, and there are available IP addresses for the subnet. To ensure that the replication instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Destination Database Access	<ul style="list-style-type: none">- Read-only During migration, the destination database is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.- Read/Write During the migration, the destination database can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks. Set the destination database to Read/Write only when you need to modify other data in the database during the migration. <p>The task cannot be modified after being created.</p>

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

- Enterprise Projects and Tags

Table 2-49 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 On the **Configure Source and Destination Databases** page, wait until the replication instance is created. Then, specify source and destination database information and click **Test Connection** for both the source and destination databases to check whether they have been connected to the replication instance. After the connection tests are successful, select the check box before the agreement and click **Next**.

 **NOTE**

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

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

Table 2-50 Self-build on ECS - source database information

Parameter	Description
Source Database Type	Select Self-built on ECS .
VPC	A dedicated virtual network in which the source database is located. It isolates networks for different services. You can select an existing VPC or create a VPC.
Subnet	A subnet provides dedicated network resources that are isolated from other networks, improving network security. The subnet must be in the AZ where the source database resides. You need to enable DHCP for creating the source database subnet.
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 – 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none">- The maximum size of a single certificate file that can be uploaded is 500 KB.- If SSL is not enabled, your data may be at risk.

 **NOTE**

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

- Scenario 2: RDS DB instance - source database configuration

Table 2-51 RDS DB instance - source database information

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

- Destination database configuration

Table 2-52 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during migration task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
Migrate Definer to User	<ul style="list-style-type: none">- Yes The Definers of all source database objects will be migrated to the user. Other users do not have permissions for database objects unless these users are authorized. For details on authorization, see How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration?- No The Definers of all source database objects will not be changed. You need to migrate all accounts and permissions of the source database in the next step.


 **NOTE**

The database username and password are encrypted and stored in the system and will be cleared after the task is deleted.

Step 4 On the **Set Task** page, select the accounts and objects to be migrated, and click **Next**.

Table 2-53 Migration types and objects

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The migration speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none">• Yes If you need to migrate accounts, see Migrating Accounts.• No During migration, accounts, permissions, and passwords are not migrated.

Parameter	Description
Filter DROP DATABASE	<p>To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations.</p> <ul style="list-style-type: none">• If you select Yes, any database deletion operations performed on the source database are not migrated during data migration.• If you select No, related operations are migrated to the destination database during data migration.
Migrate Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can choose to migrate all objects, tables, or databases based on your service requirements.</p> <ul style="list-style-type: none">• All: All objects in the source database are migrated to the destination database. After the migration, the object names will remain the same as those in the source database and cannot be modified.• Tables: The selected table-level objects will be migrated.• Databases: The selected database-level objects will be migrated. <p>If the source database is changed, click  in the upper right corner before selecting migration objects to ensure that the objects to be selected are from the changed source database.</p> <p>NOTE</p> <ul style="list-style-type: none">• If you choose not to migrate all of the databases, the migration may fail because the objects, such as stored procedures and views, in the databases to be migrated may have dependencies on other objects that are not migrated. To prevent migration failure, migrate all of the databases.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected migration object cannot contain spaces.• To quickly select the desired database objects, you can use the search function.

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

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

For details about how to handle check items that fail to pass the pre-check, see [Solutions to Failed Check Items](#).

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

 **NOTE**

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

Step 6 Compare source and destination parameters.

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

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

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

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

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

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

 **NOTE**

You can also manually enter parameter values.

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

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

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

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


- iii. Click **Next**.

Step 7 On the displayed page, specify **Start Time** and confirm that the configured information is correct and click **Submit** to submit the task.

Table 2-54 Task startup settings

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

Step 8 After the task is submitted, view and manage it on the **Online Migration Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

2.3.2 Querying the Migration Progress

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

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

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

Prerequisites

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

Procedure

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

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

- View the migration progress of structures, data, and indexes.

When a full migration is complete, the progress of each item reaches 100%.

For a full plus incremental migration, you can view the delay of the incremental migration on the **Migration Progress** page.

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

 **NOTE**

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

Transactions are synchronized as follows:

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

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

 **CAUTION**

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

- View the migration task progress. In the **Migration Details** area, locate the target migration object and click **View Details** in the **Operation** column to view the migration progress. After the incremental migration starts, the progress is not displayed. You can click the **Migration Comparison** tab to compare the data consistency.

----End

2.3.3 Viewing Migration Logs

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

Prerequisites

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

Procedure

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

Step 2 On the **Migration Logs** tab, view logs of the migration task by level.


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

----End

2.3.4 Comparing Migration Items

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

Figure 2-2 Comparison process



Comparison Scenarios

You can compare migration objects with different dimensions:

- **Object-level comparison:** It helps you compare databases, indexes, tables, views, stored procedures and functions, and sorting rules of tables. You are advised to perform the comparison after a full migration is complete.
- **Data-level comparison** is classified into row comparison and value comparison.
 - **Row comparison:** It helps you compare the number of rows in the tables to be migrated. This comparison method is recommended because it is fast.
 - **Value comparison:** It helps you check whether data in the migrated table is consistent. The comparison process is relatively slow.

When you check data consistency, compare the number of rows first. If the number of rows are inconsistent, you can then compare the data in the table to determine the inconsistent data.

Comparison Restrictions

- A comparison task can be created only when the task is in the incremental phase. When a full task is complete, DRS automatically creates object-level and row comparison tasks.
- If DDL operations were performed on the source database, you need to compare the objects again to ensure the accuracy of the comparison results.
- If data in the destination database is modified separately, the comparison results may be inconsistent.
- Currently, only tables with primary keys support value comparison. For tables that do not support value comparison, you can compare rows. Therefore, you can compare data by row or value based on scenarios.

- To prevent resources from being occupied for a long time, DRS limits the row comparison duration. If the row comparison duration exceeds the threshold, the row comparison task stops automatically. If the source database is a relational database, the row comparison duration is 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration is 30 minutes.

Impact on Databases

- Object comparison: System tables of the source and destination databases are queried, occupying about 10 sessions. The database is not affected. However, if there are a large number of objects (for example, hundreds of thousands of tables), the database may be overloaded.
- Row comparison: The number of rows in the source and destination databases is queried, which occupies about 10 sessions. The SELECT COUNT statement does not affect the database. However, if a table contains a large amount of data (hundreds of millions of records), the database will be overloaded and the query results will be returned slowly.
- Value comparison: All data in the source and destination databases is queried, and each field is compared. The query pressure on the database leads to high I/O. The query speed is limited by the I/O and network bandwidth of the source and destination databases. Value comparison occupies one or two CPUs, and about 10 sessions.

Estimated Comparison Duration

- Object comparison: Generally, the comparison results are returned within several minutes based on the query performance of the source database. If the amount of data is large, the comparison may take dozens of minutes.
- Row comparison: The SELECT COUNT method is used. The query speed depends on the database performance.
- Value comparison: If the database workload is not heavy and the network is normal, the comparison speed is about 5 MB/s.

Prerequisites

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

Creating a comparison task

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

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


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

1. Check the integrity of the database object.

Click **Validate Objects**. On the **Object-Level Comparison** tab, view the comparison result of each comparison item.

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

2. After the check is complete, compare the number of rows and values.
 - a. In the **Before You Start** pane, click **Validate All Rows/Values**.
 - b. In the displayed **Create Comparison Task** dialog box, specify **Comparison Type**, **Comparison Time**, and **Object**. Then, click **OK**.
 - **Comparison Type**: compares rows and values.
 - **Comparison Time**: You can select **Start upon task creation** or **Start at a specified time**. There is a slight difference in time between the source and destination databases during synchronization. Data inconsistency may occur. You are advised to compare migration items during off-peak hours for more accurate results.
 - **Object**: You can select objects to be compared based on the scenarios.

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

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

NOTE

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

3. Perform a double check before the cutover.

Click **Double Check During Cutover**. In the displayed **Create Comparison Task** dialog box, specify **Comparison Type**, **Comparison Time**, and **Object**. Then, click **OK**.

For details about how to view comparison details, see [Step 2.2](#).

4. Stop the migration task.

After the service system is successfully migrated to the destination database, stop the migration task to prevent operations in the source database from being synchronized to the destination database to overwrite the data. This operation only deletes the replication instance, and the migration task is still in the task list. You can view or delete the task.

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

----End

Quick Comparison

To accelerate and simplify the migration process, DRS provides the quick comparison function. You can directly perform a comparison on the migration task

list. This function can be used to compare all migration objects only when incremental migration tasks are in progress.

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

Step 2 On the **Create Comparison Task** page, select **Start upon task creation** or **Start at a specified time** and click **Yes** to start the comparison task.

----End

Viewing a Comparison Task

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

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

----End

2.3.5 Managing Objects

2.3.5.1 Migrating Accounts

Scenarios

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

MySQL Databases Operations

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

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

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

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

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

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

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

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

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

Step 3 Migrate account passwords.

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

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

Method 1: Migrate the password.

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

Method 2: Reset the password.

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

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

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

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

 **NOTE**

- The new password you set must meet the password policy of the destination database. For details, see [How Do I Change the Destination Database Password to Meet the Password Policy?](#)

----End

MongoDB Database Operations

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

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

The account information consists of account name and role.

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

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

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

----End

2.3.5.2 Parameters for Comparison

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

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

MySQL 5.6

Table 2-55 MySQL 5.6 parameters to be compared

Parameter	Type	Restart Required
connect_timeout	Common parameter	No
event_scheduler	Common parameter	No
innodb_lock_wait_timeout	Common parameter	No
max_connections	Common parameter	No
net_read_timeout	Common parameter	No
net_write_timeout	Common parameter	No
explicit_defaults_for_timestamp	Common parameter	Yes
innodb_flush_log_at_trx_commit	Common parameter	No
max_allowed_packet	Common parameter	No

Parameter	Type	Restart Required
tx_isolation	Common parameter	No
character_set_client	Common parameter	No
character_set_connection	Common parameter	No
collation_connection	Common parameter	No
character_set_results	Common parameter	No
collation_server	Common parameter	No
binlog_cache_size	Performance parameter	No
binlog_stmt_cache_size	Performance parameter	No
bulk_insert_buffer_size	Performance parameter	No
innodb_buffer_pool_size	Performance parameter	Yes
key_buffer_size	Performance parameter	No
long_query_time	Performance parameter	No
query_cache_type	Performance parameter	Yes
read_buffer_size	Performance parameter	No
read_rnd_buffer_size	Performance parameter	No
sort_buffer_size	Performance parameter	No
sync_binlog	Performance parameter	No

MySQL 5.7

Table 2-56 MySQL 5.7 parameters to be compared

Parameter	Type	Restart Required
connect_timeout	Common parameter	No
event_scheduler	Common parameter	No
innodb_lock_wait_timeout	Common parameter	No
max_connections	Common parameter	No
net_read_timeout	Common parameter	No
net_write_timeout	Common parameter	No
explicit_defaults_for_timestamp	Common parameter	No

Parameter	Type	Restart Required
innodb_flush_log_at_trx_commit	Common parameter	No
max_allowed_packet	Common parameter	No
tx_isolation	Common parameter	No
character_set_client	Common parameter	No
character_set_connection	Common parameter	No
collation_connection	Common parameter	No
character_set_results	Common parameter	No
collation_server	Common parameter	No
binlog_cache_size	Performance parameter	No
binlog_stmt_cache_size	Performance parameter	No
bulk_insert_buffer_size	Performance parameter	No
innodb_buffer_pool_size	Performance parameter	No
key_buffer_size	Performance parameter	No
long_query_time	Performance parameter	No
query_cache_type	Performance parameter	No
read_buffer_size	Performance parameter	No
read_rnd_buffer_size	Performance parameter	No
sort_buffer_size	Performance parameter	No
sync_binlog	Performance parameter	No

 NOTE

2.3.6 Task Life Cycle

2.3.6.1 Viewing Task Details

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

Prerequisites

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

Procedure

NOTE

In the task list, only tasks created by the current login user are displayed. Tasks created by different users of the same tenant are not displayed.

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

Step 2 On the displayed **Basic Information** tab, view details about the migration task. You can view information about the task, replication instance, and migration.

----End

2.3.6.2 Editing Migration Task Information

After a migration task is created, you can modify task information to identify different tasks.

The following task information can be edited:

- Task name
- Description
- Task start time

Prerequisites

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

Procedure

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

Step 2 On the **Basic Information** tab, locate the information to be modified in the **Task Information** area.




- You can click  to modify the task name and description.
 - To submit the change, click .
 - To cancel the change, click .

Table 2-57 Task information

Task Information	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).

Task Information	Description
Description	The description consists of a maximum of 256 characters and cannot contain the following special characters: !<>&'\\"

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

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

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

----End

2.3.6.3 Modifying Connection Information

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

You can modify the following information:

- Source database password
- Destination database password

NOTE

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

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click **Modify Connection Details** in the **Connection Information** area.
- Step 3** In the displayed dialog box, change the passwords of the source and destination databases and click **OK**.

----End

2.3.6.4 Modifying the Flow Control Mode

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

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

 **NOTE**

- Flow control mode takes effect only during a full migration.
- After the traffic rate is modified in the incremental migration phase, the modification takes effect when the task enters the full migration phase again.

Prerequisites

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

Method 1

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

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

----End

Method 2

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

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

----End

2.3.6.5 Editing a Migration Task

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

- Creating
- Configuration

 **NOTE**

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

Prerequisites

- You have logged in to the DRS console.

- A migration task has been created.


Method 1

- Step 1** In the task list on the **Online Migration Management** page, locate the target task and click **Edit** in the **Operation** column.
- Step 2** On the **Configure Source and Destination Databases** page, enter information about the source and destination databases and click **Next**.
- Step 3** On the **Set Task** page, select the accounts and objects to be migrated, and click **Next**.

Table 2-58 Migration types and objects

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The migration speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- Flow control mode takes effect only during a full migration.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Migrate Account	<p>During a database migration, accounts need to be migrated separately.</p> <p>There are accounts that can be migrated completely, accounts whose permissions need to be reduced, and accounts that cannot be migrated. You can choose whether to migrate the accounts based on service requirements. If you select Yes, you can select the accounts to be migrated as required.</p> <ul style="list-style-type: none">• Yes If you need to migrate accounts, see Migrating Accounts.• No During migration, accounts, permissions, and passwords are not migrated.
Filter DROP DATABASE	<p>To reduce the risks involved in data migration, DDL operations can be filtered out. You can choose not to synchronize certain DDL operations.</p> <ul style="list-style-type: none">• If you select Yes, any database deletion operations performed on the source database are not migrated during data migration.• If you select No, related operations are migrated to the destination database during data migration.

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

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

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

For details about how to handle check items that fail to pass the pre-check, see [Solutions to Failed Check Items](#).

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

 **NOTE**


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

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

 **NOTE**

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

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

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper-right corner to view the latest task status.

----End

Method 2

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

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

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

----End

2.3.6.6 Resuming a Migration Task

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

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

- Migration failed
- Paused

 **NOTE**

- If a migration task fails due to non-network problems, the system will automatically resume the task three times by default. If the failure persists, you can resume the task manually.

Prerequisites

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

Method 1

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

Method 2

- Step 1** In the task list on the **Online Migration Management** page, locate and click the task.
- Step 2** On the displayed page, click the **Migration Progress** tab, and click **Resume** in the upper left corner.
- End

2.3.6.7 Resetting a Migration Task

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

You can reset failed migration tasks in any of the following statuses:

- Migration failure status

Prerequisites

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

Method 1

- Step 1** In the task list on the **Online Migration Management** page, locate the target task and click **Reset** in the **Operation** column.
- Step 2** In the displayed dialog box, check the migration task again.
- Step 3** After the check is complete and the check success rate is 100%, click **Start** to submit the migration task again.
- End

Method 2

- Step 1** On the **Data Migration Management** page, click the target task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click the **Migration Progress** tab, and click **Reset** in the upper left corner.
- Step 3** Perform [Step 2](#) to [Step 3](#) from method 1.
- End

2.3.6.8 Pausing a Migration Task

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

You can pause the following migration tasks:

- To the cloud
 - MySQL->MySQL
- From the cloud
 - MySQL->MySQL

Prerequisites

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

Pausing a Task

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

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

NOTE

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

----End

2.3.6.9 Stopping a Migration Task

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

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

- Creating
- Configuration
- Pending start
- Full migration
- Full migration failed
- Incremental migration
- Incremental migration failed
- Paused
- Fault rectification

NOTICE

- You are advised to stop the task before performing other operations, such as disconnecting the network between the source database and the replication instance. Otherwise, an alarm indicating that the source database cannot be connected will be generated.
- For a task in the **Configuration** state, it cannot be stopped if it fails to be configured.
- For a task in the **Fault rectification** state, it cannot be stopped if the fault is being rectified.
- After a task is stopped, it cannot be resumed.

Prerequisites

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

Stopping a Task

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

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

NOTE

- Generally, triggers and events will be synchronized when you stop the task.
- If the task status is abnormal (for example, the task fails or the network is abnormal), DRS will select **Forcibly stop task** to preferentially stop the task to reduce the waiting time.
- Forcibly stopping a task will release DRS resources and will not migrate triggers and events. You have to manually migrate triggers and events.
- If you need to migrate triggers and events, restore the DRS task first. After the task status becomes normal, you can stop the task.

----End

2.3.6.10 Deleting a Migration Task

This section describes how to delete a migration task that has been completed or has failed. Deleted tasks will no longer be displayed in the task list. Exercise caution when performing this operation.

Prerequisites

- You have logged in to the DRS console.
- A migration task that has been completed or fails to be configured exists.

Deleting a Task

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

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

----End

2.3.6.11 Task Statuses

Migration statuses indicate different migration phases.

Table 2-59 lists statuses and descriptions of online migration tasks.

Table 2-59 Task status and description

Status	Description
Creating	A replication instance is being created for DRS.
Task creation failed.	Failed to create a replication instance for real-time migration.
Configuration	A replication instance is created, but the migration task is not started. You can continue to configure the task.
Pending start	The scheduled migration task has been delivered to the replication instance, waiting for the replication instance to start the migration task.
Starting	A migration task is being started.
Start failed	Failed to start a real-time migration task.
Full migration	A full migration task is being performed.
Full migration failed	Failed to perform a full migration task.
Incremental migration	An incremental migration task is being performed.
Incremental migration failed	Failed to perform an incremental migration task.
Fault rectification	A replication instance is faulty and the system automatically restores the migration task.
Paused	A real-time migration task is paused.
Stopping	The replication instance and resources used for executing the migration task are being released.
Completing	A replication instance and resources are being released.
Stopping task failed	Failed to release the replication instance and resources used by the migration task.
Completed	The task is completed and the replication instance is released.

 NOTE

- If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.
- Deleted migration tasks are not displayed in the status list.

2.4 Tag Management

Scenarios

Tag Management Service (TMS) enables you to use tags on the management console to manage resources. TMS works with other cloud services to manage tags. TMS manages tags globally, and other cloud services manage their own tags. If you have to manage a large number of tasks, you can use different tags to identify and search for tasks.

- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- Each DB instance can have up to 10 tags.

Adding a Tag

- Step 1** On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click the **Tags** tab.
- Step 3** On the **Tags** tab, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.
- Step 4** After a tag has been added, you can view and manage it on the **Tags** page.

----End

Editing a Tag

- Step 1** On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click the **Tags** tab.
- Step 3** On the **Tags** page, click **Add/Edit Tags**. In the displayed dialog box, modify the tag and click **OK**.

----End

Delete a Tag

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

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

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

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

----End

2.5 Interconnecting with Cloud Eye

2.5.1 Supported Metrics

Description

This section describes metrics reported by the Data Replication Service (DRS) to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated for DRS.

Namespace

SYS.DRS

DB Instance Monitoring Metrics

[Table 2-60](#) lists the DRS performance metrics.

Table 2-60 DRS metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
cpu_util	CPU Usage	CPU usage of the monitored object	0-100 %	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

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

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

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
apply_rows_rate	Rows Written into Destination Database per Second	Number of rows that are written to the destination database per second Unit: rows/s.	≥ 0 row/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_transactions_rate	DML TPS	Number of DML transactions written to the destination database per second.	≥ 0 transaction/s	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_ddls_number	DDL TPS	Total number of DDL transactions written into the destination database.	≥ 0 transaction	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_latency	Replication Delay	Delay (in milliseconds) of data replay.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute
apply_average_execute_time	Average Transaction Execution Time	Average execution time (RT = Execution time + Commit time) of a transaction in the destination database. The unit is millisecond.	≥ 0 ms	Monitored object: ECS Monitored instance type: replication, synchronization, and DR instances	1 minute

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

Dimensions

Key	Value
instance_id	DRS instance ID

2.5.2 Configuring Alarm Rules

Scenarios

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

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

Procedure

- Step 1** Log in to the management console.
- Step 2** Under **Management & Deployment**, click **Cloud Eye**.
- Step 3** In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
- Step 4** On the displayed **Alarm Rules** page, click **Create Alarm Rule**.

----End

2.5.3 Viewing Monitoring Metrics

Scenarios


Cloud Eye monitors the running statuses of replication, synchronization, and DR instances. You can obtain the monitoring metrics on the management console. Monitored data requires a period of time for transmission and display. The status of the monitored object displayed on the Cloud Eye page is the status obtained 5 to 10 minutes before. You can view the monitored data of a newly created DB instance 5 to 10 minutes later.

Prerequisites

An instance is running properly when in the following statuses:

- Real-time migration: Full migration and Incremental migration
- Real-time synchronization: Full synchronization and Incremental synchronization
- Real-time disaster recovery: Disaster recovery in progress

Viewing Metrics

- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and project.
- Step 3** Choose **Database > Data Replication Service**. The **Data Replication Service** page is displayed.
- Step 4** Take real-time migration as an example. On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 5 On the displayed page, click **View Metric** in the upper right corner of the page to go to the Cloud Eye console.

By default, the monitoring information about the DRS instance is displayed on this page.

Step 6 View monitoring metrics of the instance.

- On the Cloud Eye console, click the target DB instance name and click **Select Metric** in the upper right corner. In the displayed dialog box, you can select the metrics to be displayed and sort them by dragging them at desired locations.
- You can sort graphs by dragging them based on service requirements.
- Cloud Eye can monitor performance metrics from the last 1 hour, 3 hours, 12 hours, 1 day, 7 days, and 6 months.

----End

3 Backup Migration

3.1 Creating an RDS Backup Migration Task

Supported Source and Destination Databases

Table 3-1 Supported databases

Backup File	Destination DB
RDS for SQL Server full backup file versions: <ul style="list-style-type: none">• Microsoft SQL Server 2008• Microsoft SQL Server 2012• Microsoft SQL Server 2014• Microsoft SQL Server 2016• Microsoft SQL Server 2017• Microsoft SQL Server 2019	RDS for SQL Server <ul style="list-style-type: none">• Microsoft SQL Server 2008• Microsoft SQL Server 2012• Microsoft SQL Server 2014• Microsoft SQL Server 2016• Microsoft SQL Server 2017• Microsoft SQL Server 2019

Prerequisites

- You have logged in to the DRS console.
- For details about the supported DB types and versions, see [Backup Migration](#).

Precautions

This section describes constraints on backup migrations of Microsoft SQL Server databases.

Table 3-2 Precautions

Type	Restrictions
Database permissions	Before creating a backup migration task, ensure that the account has the permission to operate the RDS service.
Backup database names	<ul style="list-style-type: none">• Backup database names are case-insensitive, must be unique, and cannot be any of the following:<ul style="list-style-type: none">- msdb- master- model- tempdb- rdsadmin- resource
New database name	<ul style="list-style-type: none">• The new database name must be unique and cannot be any of the following (case-insensitive):<ul style="list-style-type: none">- msdb- master- model- tempdb- rdsadmin- resource• The new database name contains 1 to 128 characters, including letters, digits, underscores (_), and hyphens (-).
Backup file sources	<ul style="list-style-type: none">• RDS full backups: Backup files are manually or automatically created for RDS DB instances.

Type	Restrictions
Precautions	<ul style="list-style-type: none">• The available disk space of the destination database is at least 1.5 times the total data size of the backup database.• Backup database name is case-sensitive and must be the same as the database name in the backup file.• The database backup file from a database of later version cannot be restored on the instance database of an earlier version (for example, restored from version 2017 to 2016).• The restoration from Enterprise Edition to Standard Edition to Web Edition may fail. That depends on whether the features of the later version are enabled.• During a migration, if Overwrite Data is set to Yes, high availability of the destination database is disabled by default. After the migration is complete, high availability is restored automatically.• During a migration, stop writing transactions to the destination database.• If a primary/standby switchover of the destination database is performed, the backup migration fails. In this case, the migration task cannot be restored.

Procedure

This section describes how to create an RDS full backup migration task. You can use the full backups of Microsoft SQL Server DB instances in the cloud to migrate data.

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

Step 2 On the **Select Backup** page, specify information about the task and backup files. Then, click **Next**.

Table 3-3 Task information

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

Table 3-4 Backup file information

Parameter	Description
Database Type	Select Microsoft SQL Server .
Backup File Source	Select RDS full backup . NOTE Select a backup file whose status is Completed .
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">• Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.• After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

Step 3 On the **Select Destination** page, specify database information and click **Next**.

Table 3-5 Database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. If no RDS DB instance is available, you can create one. For details, see .
Backup Database Name	After you select the destination RDS DB instance, all databases to be restored are automatically displayed. You can select databases to be restored as required and rename them. <ul style="list-style-type: none">• Backup Database Name: Name of the database to be restored.• New Database Name: The backup database name must consist of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (_). If the name is not specified, the original database name is used. NOTE <ul style="list-style-type: none">• The backup database can be renamed. A maximum of 100 backup databases can be created.• The new database name cannot be the same as the name of any other database in the source.

Step 4 On the **Confirm Task** page, check configuration details, read and agree to the agreement, and click **Next**.

 **NOTE**

If the SQL Server source contains non-clustered index tables, the index information of non-clustered index tables will become invalid after the SQL Server backups are restored to a new database. For the best performance, rebuild the indexes after the backup migration. In addition, the backup files store only database-level information. If the SQL Server source contains some instance-level configurations, such as login, permission, DBlink, and job, migrate these configurations by referring to [Manual Configuration](#).

Step 5 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

3.2 Creating a Backup Using OBS Buckets

Supported Source and Destination Databases

Table 3-6 Supported databases

Backup File	Destination DB
On-premises and other cloud's Microsoft SQL Server backup file versions: <ul style="list-style-type: none">• Microsoft SQL Server 2000• Microsoft SQL Server 2005• Microsoft SQL Server 2008• Microsoft SQL Server 2012• Microsoft SQL Server 2014• Microsoft SQL Server 2016• Microsoft SQL Server 2017	RDS for SQL Server <ul style="list-style-type: none">• Microsoft SQL Server 2008• Microsoft SQL Server 2012• Microsoft SQL Server 2014• Microsoft SQL Server 2016• Microsoft SQL Server 2017

Prerequisites

- You have logged in to the DRS console.
- For details about the supported DB types and versions, see [Backup Migration](#).

Precautions

This section describes constraints on backup migrations of Microsoft SQL Server databases.

Table 3-7 Precautions

Type	Restrictions
Database permissions	Before creating a backup migration task, ensure that the current account has the OBS and bucket permissions.
Backup database names	<ul style="list-style-type: none">• Backup database names are case-insensitive, must be unique, and cannot be any of the following:<ul style="list-style-type: none">- msdb- master- model- tempdb- rdsadmin- resource• The name of the OBS bucket contains 1 to 256 characters, including letters, digits, underscores (_), and hyphens (-).
New database name	<ul style="list-style-type: none">• The new database name must be unique and cannot be any of the following (case-insensitive):<ul style="list-style-type: none">- msdb- master- model- tempdb- rdsadmin- resource• The new database name contains 1 to 128 characters, including letters, digits, underscores (_), and hyphens (-).
Local backup files	<ul style="list-style-type: none">• The suffix of the backup file to be uploaded to an OBS bucket must be .bak. Only .bak files can be selected in the OBS backup file list.• The backup file name contains 1 to 200 characters.• The backup file name can contain only letters, digits, hyphens (-), and underscores (_).• Backup files are classified into full backup files and log backup files.
Backup file sources	<ul style="list-style-type: none">• OBS buckets: store local backup files in its root directory.

Type	Restrictions
Precautions	<ul style="list-style-type: none">• The OBS bucket and DB instance must be in the same region.• The available disk space of the destination database is at least 1.5 times the total data size of the backup database.• Backup database name is case-sensitive and must be the same as the database name in the backup file.• The database backup file from a database of later version cannot be restored on the instance database of an earlier version (for example, restored from version 2017 to 2016).• The restoration from Enterprise Edition to Standard Edition to Web Edition may fail. That depends on whether the features of the later version are enabled.• During a migration, if Overwrite Data is set to Yes, high availability of the destination database is disabled by default. After the migration is complete, high availability is restored automatically.• During a migration, stop writing transactions to the destination database.• If a primary/standby switchover of the destination database is performed, the backup migration fails. In this case, the migration task cannot be restored.

Procedure

This section describes how to create a backup migration task using the backup stored in the OBS bucket on the DRS console.

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

Step 2 On the **Select Backup** page, specify information about the task and backup files. Then, click **Next**.

Table 3-8 Task information

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

Table 3-9 Backup file information

Parameter	Description
Database Type	Database type of the backup file. Select Microsoft SQL Server.
Backup File Source	Select OBS Bucket .
Bucket Name	Select a bucket and a backup file stored in the directory of the bucket. NOTE <ul style="list-style-type: none">The backup file of the Microsoft SQL Server must be named in the .bak format in the OBS bucket directory. Multiple backup files can be selected at the same time.A database cannot be split into different files for upload.
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">Tags a task. This configuration is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

Step 3 On the **Select Destination** page, specify database information and click **Next**.

Table 3-10 Microsoft SQL Server database information

Parameter	Description
Destination RDS DB Instance Name	Select a destination RDS DB instance. If no RDS DB instance is available, you can create one.
Backup File Format	Select Full or Incremental . <ul style="list-style-type: none">Full: indicates full backup files.Incremental: indicates log backup files. NOTE <ul style="list-style-type: none">Before performing an incremental restoration, you need to perform a full restoration.To migrate databases at a time, you need to stop services first and upload full backups for restoration.To ensure service continuity during a migration, perform a full backup and multiple incremental backups, and then upload the backup data to OBS for restoration.

Parameter	Description
Last Backup	<p>A complete restoration involves restoring several incremental backup files. The destination DB instance cannot be modified while it is being restored. It will not be available until the last backup file is restored. If you select Yes and restore the instance, any incremental restoration will not be performed. You can select the last backup in the following scenarios:</p> <ul style="list-style-type: none">• Perform a one-time full migration.• The selected backup file is the last one to be restored.• Select No for the full+incremental backup migration scenario where you continue to restore databases using incremental backup files after a full backup restoration is performed. During this period, the destination DB instance is being restored and in the unavailable state.
Overwrite Data	<p>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. Exercise caution when performing this operation.</p> <p>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.</p>
Perform Pre-verification	<p>Specifies whether to perform pre-verification on the backup migration task. The default value is Yes.</p> <ul style="list-style-type: none">• 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	<p>You can restore all or part of databases. All of databases are restored by default.</p> <ul style="list-style-type: none">• All: Restore all databases from the backup file. You do not need to enter the name of the database to be 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.

Parameter	Description
Reset Database Name	<p>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.</p> <p>Constraint:</p> <ul style="list-style-type: none">• 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. <p>NOTE The database name can be reset only when Backup Type is set to Full and Restore Database is set to All.</p>
Backup Database Name	<p>If Restore Database is set to Custom, you need to specify Backup Database Name.</p> <p>The backup database name must be the same as that in the backup file and consist of 1 to 256 characters. It can contain only letters (case-sensitive), digits, hyphens (-), and underscores (_).</p> <p>DRS also allows you to create an alias for the database to be restored.</p> <ul style="list-style-type: none">• If you choose to restore databases in full backup mode, you can set the alias of the database based on the site requirements. The alias is also stored in the destination database.• If databases are restored in incremental backup mode, alias is not supported. <p>NOTE The backup database can be renamed. Up to 100 backup databases can be created.</p>

Step 4 On the **Confirm Task** page, check configuration details, read and agree to the agreement, and click **Next**.

 **NOTE**

If the SQL Server source contains non-clustered index tables, the index information of non-clustered index tables will become invalid after the SQL Server backups are restored to a new database. For the best performance, rebuild the indexes after the backup migration. In addition, the backup files store only database-level information. If the SQL Server source contains some instance-level configurations, such as login, permission, DBlink, and job, migrate these configurations by referring to [Manual Configuration](#).

Step 5 In the task list on the **Backup Migration Management** page, check whether the task is in the **Restoring** status. If the migration is successful, the task status becomes **Successful**.

----End

3.3 Task Management

3.3.1 Viewing Task Details

The **Basic Information** page of a backup migration task displays information about the task, backup files, and databases. This section describes how to view details about current migration task.

Prerequisites

You have logged in to the DRS console.

Procedure

NOTE

In the task list, only tasks created by the current login user are displayed. Tasks created by different users of the same tenant are not displayed.

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

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

You can view information about the task, backup files, and databases.

----End

3.3.2 Editing Migration Task Information

After a backup migration task is created, you can modify basic information about the task to identify different tasks.

The following task information can be edited:


- Task name
- Description

Prerequisites

You have logged in to the DRS console.

Procedure

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

Step 2 On the **Basic Information** page, click  next to the information to modify.



- To submit the change, click .
- To cancel the change, click .

Table 3-11 Task information description

Task Information	Description
Task name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).
Description	The description consists of a maximum of 256 characters and cannot contain the following special characters: ! = < > ' & "

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

----End

3.3.3 Viewing Migration Logs

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

Prerequisites

You have logged in to the DRS console.

Procedure

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

Step 2 On the displayed page, click **Migration Logs** to view the logs generated during the migration.

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

----End

3.3.4 Deleting a Migration Task

This section describes how to delete a migration task. Deleted tasks will no longer be displayed in the task list. Exercise caution when performing this operation.

Prerequisites

You have logged in to the DRS console.

Deleting a Task

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

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

----End

3.3.5 Task Statuses

Migration statuses indicate different migration phases.

Table 3-12 lists statuses and descriptions of backup migration tasks.

Table 3-12 Backup migration task statuses

Status	Description
Restoring	A backup file is being restored to the destination database.
Successful	A backup file has been restored to the destination database.
Failed	A backup file fails to be restored to the destination database.
Check failed	A backup file is unavailable.

 **NOTE**

Deleted migration tasks are not displayed in the status list.

3.4 Tag Management

Scenarios

Tag Management Service (TMS) enables you to use tags on the management console to manage resources. TMS works with other cloud services to manage tags. TMS manages tags globally, and other cloud services manage their own tags. If you have to manage a large number of tasks, you can use different tags to identify and search for tasks.

- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- Each DB instance can have up to 10 tags.

Adding a Tag

- Step 1** On the **Backup Migration Management** page, click the target migration task name in the **Task Name/ID** column.
- Step 2** On the **Basic Information** tab, click the **Tags** tab.
- Step 3** On the **Tags** tab, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.
- Step 4** After a tag has been added, you can view and manage it on the **Tags** page.

----End

Editing a Tag

- Step 1** On the **Backup Migration Management** page, click the target migration task name in the **Task Name/ID** column.
 - Step 2** On the **Basic Information** tab, click the **Tags** tab.
 - Step 3** On the **Tags** page, click **Add/Edit Tags**. In the displayed dialog box, modify the tag and click **OK**.
- End

Delete a Tag

- Step 1** On the **Backup Migration Management** page, click the target migration task name in the **Task Name/ID** column.
 - Step 2** On the **Basic Information** tab, click the **Tags** tab.
 - Step 3** On the **Tags** page, locate the tag to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.
 - Step 4** After the tag is deleted, it will no longer be displayed on the **Tags** page.
- End

4 Real-Time Synchronization


4.1 Data Synchronization Topologies

DRS real-time synchronization supports multiple topology types. You can plan the topology types as required. For details, see the following content.

 **NOTE**


To ensure data consistency, do not modify the synchronization objects in the destination database.

One-to-One Real-Time Synchronization




You can create a one-to-one synchronization task.

One-to-Many Real-Time Synchronization



You need to create multiple synchronization tasks to implement one-to-many real-time synchronization. For example, to synchronize data from instance A to instances B, C, and D, you need to create three synchronization tasks.

Many-to-One Real-Time Synchronization



You need to create multiple synchronization tasks to implement many-to-one real-time synchronization. For example, to synchronize data from instances B, C, and D to instance A, you need to create three synchronization tasks. Multiple tables can be synchronized to one table.

4.2 To the Cloud

4.2.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 4-1 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none"> On-premises MySQL databases MySQL databases on an ECS MySQL databases on other clouds RDS for MySQL 	<ul style="list-style-type: none"> RDS for MySQL

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - Start your synchronization 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. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - Data-Level Comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

- For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-2 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none">• The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT• The destination database user must have the following permissions: The root account of RDS for MySQL has the following permissions by default: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none"> ● Only tables, primary key indexes, unique indexes, common indexes, store procedures, views, and functions can be synchronized. ● Table names cannot be mapped for tables on which views, stored procedures, and functions depend. ● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized. ● During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, resulting in inconsistent objects. ● If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. ● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized. ● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent. <ul style="list-style-type: none"> - If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out. - If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out. - You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent.

Type	Restrictions
Source database	<ul style="list-style-type: none">● The source database names cannot contain non-ASCII characters, or the following characters: '<`>/\''● The source table and view names cannot contain non-ASCII characters, or the following characters: '<`>/\''● The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1.● During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format.● If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail.<ul style="list-style-type: none">- If the source database is an on-premises MySQL database, set expire_logs_days to specify the binlog retention period. Set expire_logs_day to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.- If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide.● GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.● During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• Data cannot be synchronized from a newer version database to an older version database.• The destination DB instance is running properly. If the destination DB instance type is primary/standby, the replication status must also be normal.• The destination DB instance must have sufficient storage space.• The character set of the destination database must be the same as that of the source database.• The time zone of the destination database must be the same as that of the source database.• If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.• During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.• If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.• The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (').

Type	Restrictions
Precautions	<ul style="list-style-type: none">● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys● Only MySQL to MySQL synchronization supports many-to-one synchronization. During table-level many-to-one synchronization, tables without primary keys cannot exist in the source database.● If you create many-to-one synchronization tasks, the system automatically creates a parent task to associate multiple synchronization tasks after the tasks are started. The parent task is named in the <i>DRS-Group-Destination DB instance name</i> format.● If the sources and destinations are RDS instances, database mapping is required.● The source and destination databases cannot contain tables that have the same names but do not have primary keys.● The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.● If the source and destination sides are RDS MySQL instances, tables encrypted using TDE cannot be synchronized.● If the source MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key.● The destination table can contain more columns than the source table. However, the following failures must be avoided:<ul style="list-style-type: none">– Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.– Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.● The source database does not support point-in-time recovery (PITR).

Type	Restrictions
	<ul style="list-style-type: none">● The destination database cannot be restored to a point in time when a full synchronization was being performed.● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails.● The partitioned table does not support column mapping.● Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.● After a task is created, the destination database cannot be set to read-only.● A real-time synchronization task may fail due to the change of the username and password of the source or destination database. If it happens, rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization.● If the source or destination database port is changed during data synchronization, the synchronization task fails. You can rectify the fault as follows:<ul style="list-style-type: none">- If the source database port is wrong, correct the port number on the DRS console and then retry the synchronization task.- If the destination database port is wrong, DRS automatically changes the port to the correct one, and then you need to retry the synchronization task. Generally, do not modify the port number during synchronization.● During data synchronization, if the source database is on an RDS instance that does not belong the current cloud platform, the IP address cannot be changed. If the source database is on an RDS DB instance on the current cloud platform, the system automatically changes the IP address to the correct one. Then, retry the task to continue the synchronization. Therefore, changing the IP address is not recommended.● To ensure data consistency, do not perform operations (including but not limited to DDL and DML operations) on the destination database during the synchronization.● Data inconsistency may occur when the MyISAM table is modified during a full synchronization.● DDL statements can be synchronized during incremental synchronization.● Incremental synchronization supports table renaming. The source and destination tables must be selected in the

Type	Restrictions
	<p>synchronization object list. Full synchronization does not support the RENAME command.</p> <ul style="list-style-type: none">You can add additional objects during an incremental synchronization.

Procedure

This section describes how to synchronize data from a MySQL database to an RDS MySQL database. To configure other storage engines, you can refer to the following procedures.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Table 4-3 Task and recipient description

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

- Synchronization instance information

Table 4-4 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database is a database in the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect

Parameter	Description
Destination DB Instance	The RDS DB instance you created. NOTE <ul style="list-style-type: none">- The destination DB instance cannot be a read replica.- The source and destination DB instances can be the same DB instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.
Synchronization Mode	Available options: Full+Incremental and Incremental <ul style="list-style-type: none">- Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.- Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

- Enterprise Projects and Tags

Table 4-5 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .

Parameter	Description
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, 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**.

- Source database information

Table 4-6 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 - 65535
Database Username	The username for accessing the source database.
Database Password	<p>The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:</p> <p>If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details. In the displayed dialog box, change the password.</p>
SSL Connection	<p>SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate.</p> <p>NOTE</p> <ul style="list-style-type: none"> - The maximum size of a single certificate file that can be uploaded is 500 KB. - If SSL is not enabled, your data may be at risk.

 NOTE

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

- Destination database information

Table 4-7 Destination database settings

Parameter	Description
DB Instance Name	The RDS DB instance selected during synchronization task creation. This parameter cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.


 NOTE

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

- Step 4** On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Table 4-8 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- The flow control mode takes effect only in the full synchronization phase.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">• 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. <p>If data conflicts occur, you can select Ignore, Overwrite, or Report error in any of the following scenarios.</p> <ul style="list-style-type: none">• Data exists in the destination database.• Multiple source databases are synchronized to one destination database.• Data in the destination database is updated manually.

Parameter	Description
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none">• 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.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize data based on service requirements.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none">• If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names.<ul style="list-style-type: none">– If the database table name contains characters other than letters, digits, and underscores (<code>_</code>), or the mapped database table name contains hyphens (<code>-</code>) and number signs (<code>#</code>), the name length cannot exceed 42 characters. <p>NOTE</p> <ul style="list-style-type: none">• To quickly select the desired database objects, you can use the search function.• If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

- If data processing is not required, click **Next**.
- If data processing is required, select **Data filtering**, **Additional Column**, or **Processing Columns**. For details about how to configure related rules, see **Processing Data**.

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


 **NOTE**

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

- Step 7** On the displayed page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

Table 4-9 Task startup settings

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

- Step 8** After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.
- You can view the task status. For more information about task status, see [Task Statuses](#).
 - You can click  in the upper-right corner to view the latest task status.
 - By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.2.2 From PostgreSQL to PostgreSQL

Supported Source and Destination Databases

Table 4-10 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12 and 13)• ECS database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12 and 13)• Other cloud database (PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12 and 13)• RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13)	RDS for PostgreSQL (9.5, 9.6, 10, 11, 12, 13) NOTE The major version of the destination database must be the same as or later than that of the source database.

Supported Synchronization Objects

Table 4-11 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-11 Supported synchronization objects

Type	Notes
Objects	<ul style="list-style-type: none">• Instance-level synchronization is not supported. Only one database can be synchronized at a time. Multiple DRS tasks are required to synchronize multiple databases.• Supported field types: Digit, currency, character, binary, date/time, boolean, enumeration, geometry, network address, bit, text search, UUID, XML, JSON, array, compound, and range.• Scope of full synchronization<ul style="list-style-type: none">- Schema, table, index, constraint, view, materialized view, sequence, stored procedure, rule, trigger, foreign key, sorting rule, plug-in, code conversion information, aggregate function, operator, statistics extension, conversion information, text search configuration, function, data type, type conversion, user, and event trigger, text search parsers, and text search templates During the table-level synchronization, only tables, views, materialized views, sequences, and users can be synchronized.- Objects that are not supported: system schemas (schemas starting with pg_, information_schema, sys, utl_raw, dbms_lob, dbms_output, and dbms_random), system catalogs, system users, tablespaces, foreign-data wrappers, foreign servers, user mappings, publications, and subscriptions

Type	Notes
	<p>NOTE</p> <p>The restrictions on the objects that can be synchronized are as follows:</p> <ul style="list-style-type: none">• Object name: The database name cannot contain "+" %?\<>, the schema name and table name cannot contain ".\<>, and the column name cannot contain double quotation marks (") and single quotation marks (').• Table: Temporary tables are not synchronized. During table-level synchronization, table constraints, indexes, and rules are synchronized, except for table triggers.• Schema: Permissions of the public schema are not synchronized. During table-level synchronization, the permissions of existing schemas in the destination database are synchronized.• Function: Do not synchronize C language functions or functions with the leakproof or support attribute.• Plug-in: The metadata of plug-ins is not synchronized.• Data type: Basic data types are not synchronized.• Type conversion: The binary coercion type cannot be converted.• Event trigger: Event triggers can be synchronized only when the destination database version is RDS for PostgreSQL 11.11 or later.• Text search parser: Text search parsers can be synchronized only when the destination database version is RDS for PostgreSQL 11.11 or later.• Text search template: Text search templates can be synchronized only when the destination database version is RDS for PostgreSQL 11.11 or later.• User: Existing users in the destination database, superuser, replication, and bypassrsls attributes of users, and member relationships of superuser users are not synchronized. If the object owner or grantor is superuser, its owner or grantor is not synchronized. If the destination database is RDS for PostgreSQL DB instance, the password of the user to be synchronized cannot contain the username. During table-level synchronization, the default access permissions of source database users are not synchronized. <ul style="list-style-type: none">• Scope of incremental synchronization<ul style="list-style-type: none">- Some DML statements, including INSERT, UPDATE, and DELETE, can be synchronized.- Some DDL operations can be synchronized, including TRUNCATE (PostgreSQL 11 or later), CREATE SCHEMA, CREATE TABLE, DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, and ALTER VIEW. <p>During table-level synchronization, only the following DDL operations can be synchronized: TRUNCATE (PostgreSQL 11 or later), DROP TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME).</p>

Type	Notes
	<ul style="list-style-type: none">- Not synchronized: DML statements of unlogged tables and temporary tables <p>NOTE The source database captures DDL statements using event triggers and records them in specific tables, so you need to create event triggers and functions in the source database in advance. For details, see Creating Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</p>

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-12](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

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

Table 4-12 Database account permission

Type	Full	Full+Incremental
Source database user	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)	<p>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</p> <p>NOTE</p> <ul style="list-style-type: none"> • 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: <ol style="list-style-type: none"> 1. Add host replication <code><src_user_name></code> <code><drs_instance_ip>/32 md5</code> before all configurations in the pg_hba.conf file of 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.

Type	Full	Full+Incremental
Destination database user	<ul style="list-style-type: none">● Database-level: The CREATEDB permission is required.● Table-level:<ul style="list-style-type: none">- To synchronize databases, the CREATEDB permission is required.- To synchronize a schema, the CONNECT and CREATE permissions for the database that contains the schema are required.- To synchronize objects in a schema, the CONNECT permission for the database that contains the schema, and the USAGE and CREATE permissions for the schema that contain the object are required.● Synchronization user: The CREATEROLE permission is required.● Synchronization user permissions: The default privilege cannot be modified. Otherwise, the object permissions of the destination database may be inconsistent with those of the source database. <p>NOTE To synchronize event triggers, text search parsers, and text search templates, the destination database version must be RDS for PostgreSQL 11.11 or later, and the destination database user must be user root or a member of user root.</p>	

Suggestions

⚠ CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - Start your synchronization 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.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.

- The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
- When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
- If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
- Data-Level Comparison
To obtain accurate comparison results, start data comparison at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

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

Table 4-13 Precautions

Type	Constraints
Starting a task	<ul style="list-style-type: none">● Source database parameter requirements:<ul style="list-style-type: none">- The partition table trigger of the source database cannot be set to disable.- For a full synchronization task, the source database can be a standby database, but hot_standby_feedback must be set to on. For an incremental synchronization task, the source database cannot be a standby database.- To perform incremental synchronization: The pg_hba.conf file of the source database contains the following configuration: <pre>host replication all 0.0.0.0/0 md5</pre> The wal_level value of the source database must be logical. The test_decoding plug-in has been installed on the source database. The replica identity attribute of tables that do not have primary keys in the source database must be full. The max_replication_slots value of the source database must be greater than the number of used replication slots. The max_wal_senders value of the source database must be greater than or equal to the max_replication_slots value. If the toast attribute of the primary key column in the source database is main, external, or extended, the replica identity attribute must be full.● Source database object requirements:<ul style="list-style-type: none">- Triggers with the same name cannot exist in the source database.- The objects that have dependencies must be synchronized at the same time. Otherwise, the synchronization may fail.● Destination database parameter requirements:<ul style="list-style-type: none">- The block_size value of the destination database must be greater than or equal to that of the source database.- The lc_monetary values of the source and destination databases must be the same.- To perform incremental synchronization, if the synchronization object contains foreign keys, triggers, or event triggers, set session_replication_role of the destination database to replica. After the synchronization is complete, change the value to origin.● Destination database object requirements:<ul style="list-style-type: none">- The destination database cannot contain objects with the same type and name as the objects to be synchronized,

Type	Constraints
	<p>including databases, schemas, and tables. System databases, system schemas, and system tables are excluded.</p> <ul style="list-style-type: none"> - The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail. <p>● Other notes:</p> <ul style="list-style-type: none"> - When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value +original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, the original constraint name on the table is changed to c_+ hash value + original constraint name (which may be truncated) + _key. - Before starting a full+incremental or incremental synchronization task, ensure that no long transaction is started in the source database. Starting the long transaction will block the creation of the logical replication slot and cause the task to fail. - For a full+incremental or incremental synchronization task, if an internal error occurs during the pre-check and the task stops before it is started, check and delete the streaming replication slot by referring to Forcibly Stopping Synchronization of PostgreSQL to avoid residual streaming replication slots in the source database. - If you choose to synchronize DDL statements, ensure that the DDL statements executed on the source database are compatible with the destination database.

Type	Constraints
	<p>NOTE DDL statements are captured using event triggers in the source database, recorded in a specific table, and then synchronized to the destination database. You need to create event triggers and functions in the source database before starting a task. For details, see Creating Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL.</p>
Full synchronization	<ul style="list-style-type: none">• Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail.• Do not run any DDL statement in the source database. Otherwise, data may be inconsistent or the task may fail.• Do not write data to the destination database. Otherwise, data may be inconsistent.
Incremental synchronization	<ul style="list-style-type: none">• Do not change the port of the source and destination databases, or change or delete the passwords and permissions of the source and destination database users. Otherwise, the task may fail.• Do not change the primary key or unique key (if the primary key does not exist) of the source database table. Otherwise, incremental data may be inconsistent or the task may fail.• Do not modify the replica identity attribute of tables in the source database. Otherwise, incremental data may be inconsistent or the task may fail.• Do not write data to the destination database. Otherwise, data may be inconsistent.• During database-level synchronization, if a table without a primary key is added to the source database, you must set replica identity of the table to full before writing data. Otherwise, data may be inconsistent or the task may fail.• During database-level synchronization, if a primary key table is added to the source database and the toast attribute of the primary key column is main, external, or extended, the replica identity attribute of the table must be set to full before writing data. Otherwise, data may be inconsistent or the task may fail.
Synchronization comparison	<ul style="list-style-type: none">• You are advised to compare data in the source database during off-peak hours to prevent inconsistent data from being falsely reported and reduce the impact on the source database and DRS tasks.• During incremental synchronization, if data is written to the source database, the comparison results may be inconsistent.• Data cannot be compared during full synchronization.• Do not limit the synchronization speed during data comparison.

Type	Constraints
Stopping a task	<ul style="list-style-type: none">● Stop a task normally:<ul style="list-style-type: none">– The destination database sequence value is automatically reset. The auto-increment sequence value is the source database sequence value plus the security margin, and the auto-decrement sequence value is the source database sequence value minus the security margin. The default security margin is 10,000. If users are synchronized, the user memberships are automatically synchronized after the task is complete.– When a full+incremental synchronization task is complete, the streaming replication slot created by the task in the source database is automatically deleted.– If the value of destination database session_replication_role is replica when the full+incremental synchronization task is complete, change the value to origin.● Forcibly stop a task:<ul style="list-style-type: none">– You need to manually update the sequence value in the destination database. For details, see Forcibly Stopping Synchronization of PostgreSQL.– To forcibly stop a full+incremental real-time synchronization task, you need to manually delete the replication slots that may remain in the source database. For details, see Forcibly Stopping Synchronization of PostgreSQL.– If the value of destination database session_replication_role is replica, change it to origin to forcibly stop the full+incremental synchronization task.

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section uses synchronization from PostgreSQL to PostgreSQL as an example to describe how to configure a real-time synchronization task in the VPC network scenario.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Table 4-14 Task and recipient description

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

- Synchronization instance details

Table 4-15 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select PostgreSQL .
Destination DB Engine	Select PostgreSQL .
Network Type	Available options: VPC , Public network and VPN or Direct Connect . VPC is used as an example.
Destination DB Instance	The RDS for PostgreSQL DB instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none">- 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.

- Enterprise Projects and Tags

Table 4-16 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">- This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.- After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

- Step 3** 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**.

Establish the connectivity between the DRS instance and the source and destination databases.

- **Network connectivity:** Ensure that the source and destination databases accept connections from the IP address of the DRS instance. To access databases over a public network, configure the database to accept connections from the EIP of the DRS instance. To access databases over a VPC, VPN, or Direct Connect network, configure the database to accept connections from the private IP address of the DRS instance.

- **Account connectivity:** Ensure that the source and destination databases allows connections from the DRS instance using the username and password.

 **NOTE**

The source database can be an ECS database or an RDS instance. Configure parameters based on the database type.

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

Table 4-17 Self-build on ECS - source database information

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

 **NOTE**

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

- Scenario 2: RDS DB instance - source database configuration

Table 4-18 RDS DB instance - source database information

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

Table 4-19 Destination database settings

Parameter	Description
DB Instance Name	The RDS PostgreSQL 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.


NOTE

The username and password of the source and destination databases are encrypted and stored in the databases and the synchronization instance during the synchronization. After the task is deleted, the username and password are permanently deleted.

Step 4 On the **Set Synchronization Task** page, select the synchronization objects and accounts and click **Next**.

Table 4-20 Synchronization Object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- The flow control mode takes effect only in the full synchronization phase.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">• 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.
Synchronize	<p>Available options: Index, Incremental DDLs, and Populate materialized views during the full synchronization phase</p> <p>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.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none">• Database-level synchronization: In full synchronization, the selected databases and the inventory data of the database objects are synchronized. In incremental synchronization, the DML and some DDL statements of all tables except unlogged tables and temporary tables are synchronized.• Table-level synchronization: In full synchronization, the inventory data of the selected tables, sequences, views, or materialized views is synchronized. In incremental synchronization, the DML and some DDL statements of the selected tables are synchronized.• If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. When a schema name or table name is mapped, to prevent conflicts between indexes and constraint names, the original index name in the table is changed to the following format after synchronization: i_+hash value+original index name (which may be truncated)+_key The hash value is calculated based on the original schema name_original table name_original index name. Similarly, after the synchronization, the original constraint name on the table is changed to c_ + hash value + original constraint name (which may be truncated) + _key. <p>NOTE</p> <ul style="list-style-type: none">• To quickly select the desired database objects, you can use the search function.• If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.
Synchronize Account	<p>During the synchronization, you can synchronize accounts based on your service requirements.</p> <p>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.</p>

Step 5 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**.

NOTE


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

Step 6 On the displayed page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

Table 4-21 Task startup settings

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

Step 7 After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper-right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.2.3 From Oracle to MySQL

Supported Source and Destination Databases

Table 4-22 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises databases (Oracle 10g, 11g, 12c, 18c and 19c)• Self-built databases on ECS (Oracle 10g, 11g, 12c, 18c and 19c)	<ul style="list-style-type: none">• RDS for MySQL

Supported Synchronization Objects

Table 4-23 lists the objects that can be synchronized in different scenarios. DRS will automatically check the objects you selected before the synchronization.

Table 4-23 Supported synchronization objects

Type	Precautions
Objects	<ul style="list-style-type: none"> ● Object level: table level. ● Supported synchronization objects: <ul style="list-style-type: none"> - Databases, table structures, primary keys, unique keys, normal indexes, and table data can be synchronized. Other database objects, such as stored procedures, triggers, functions, sequences, packages, synonyms, and users, cannot be synchronized. - In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported. - In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported. - During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail. - Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by primary or foreign keys - Partitions in the table structure cannot be synchronized. Partitioned tables are changed to non-partitioned tables after being synchronized to the destination database. - Tables whose default values contain expressions of functions cannot be synchronized. - Temporary tables in the source database cannot be synchronized. - An empty source database cannot be synchronized. - If the table contains only LOB columns, data inconsistency may occur. - If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database. - For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched.

Type	Precautions
	<p>NOTE</p> <ul style="list-style-type: none">• Database object names, such as the database name and table name, support English characters and symbols such as #, \$, and _ . DRS does not support non-ASCII characters or special characters .>`<'\, ?!"• Object names will be converted to lowercase letters after being synchronized to the destination database. To avoid synchronization failures, ensure that the selected source database tables do not contain tables with the same name but different letter cases.

Database Account Permission Requirements

To start a synchronization task, the source and destination database users must meet the requirements in the following table. Different types of synchronization tasks require different permissions. For details, see [Table 4-24](#). DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

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

Table 4-24 Database account permission

Type	Full	Full+Incremental
Source database user	CREATE SESSION, SELECT ANY DICTIONARY, and SELECT (GRANT SELECT <userName.tbName> to drsUser) for single tables	<ul style="list-style-type: none"> • For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). • 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. • For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. • During incremental synchronization, enable PK, UK, 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

Type	Full	Full+Incremental
		<p>synchronization, ensure that the preceding settings are always enabled.</p> <ul style="list-style-type: none">• 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.
Destination database user	The user must have the SELECT, INSERT, CREATE, DROP, UPDATE, ALTER, DELETE and INDEX permissions.	

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - Start your synchronization 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.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - Data-Level Comparison
To obtain accurate comparison results, [compare data](#) at a specified time point during off-peak hours. If it is needed, select **Start at a specified time**

for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

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

Table 4-25 Precautions

Type	Restrictions
Starting a task	<ul style="list-style-type: none">● Source database parameter requirements:<ul style="list-style-type: none">- Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.● Source database object requirements:<ul style="list-style-type: none">- The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1.- The default values of the source database can be to_date and sys_guid functions. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed: The default value may be left empty. The table fails to be created. As a result, the object comparison result is inconsistent or the task fails.- The maximum row length of Oracle cannot exceed 8 KB, excluding BLOB and TEXT columns which will be automatically converted to the text and blob types of MySQL. The reason is that the MySQL InnoDB restricts the row length to 8 KB.- The primary key or unique key column cannot contain values of the string data type when you map the MySQL data types to the character data types in Oracle because MySQL cannot tell spaces in data. Otherwise, data inconsistency and deadlock may occur.- The values of binary_float and binary_double cannot be set to Nan, Inf, or -Inf because MySQL does not support these values. DRS converts the three values to 0 and saves them by default.- MySQL does not support the synchronization of the check constraints of Oracle.- AUTO_PK_ROW_ID cannot be used as a column name in Oracle because it is a reserved column name in MySQL 5.7 and cannot be created.- Ensure that the precision of the number(p, s) field in the Oracle database does not exceed the precision range p: [1, 38], s:[p-65, min(p, 30)]. The value of s depends on the value of p. The lower limit is p-65, and the upper limit is the minimum value of p or 30. For example, when p is 1, the value range of s is [-64, 1]. When p is 38, the value range of s is [-27, 30]. The value of the int field cannot

Type	Restrictions
	<p>exceed the precision range of (65, 0). The digit range of MySQL is smaller than that of Oracle.</p> <ul style="list-style-type: none">- The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record.- The Default User statement is not supported in MySQL. <ul style="list-style-type: none">● Destination database parameter requirements:<ul style="list-style-type: none">- During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.● Destination database object requirements:<ul style="list-style-type: none">- The time zone settings of the source and destination database must be the same.- When you select to synchronize the table structure, the destination instance cannot contain the database to be synchronized.- The storage of the destination database should be about 1.5 times greater than the storage of the source database.- If the destination database version is earlier than 5.7.7, the index column length cannot exceed 767 bytes. If the destination database version is later than 5.7.7, the length cannot exceed 3072 bytes.- Do not use foreign keys for tables during incremental synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures.- The destination table can contain more columns than the source table. However, the following failures must be avoided:<ul style="list-style-type: none">Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.● Other notes:

Type	Restrictions
	<ul style="list-style-type: none"><li data-bbox="627 297 1366 360">- If the data types are incompatible, the synchronization may fail.<li data-bbox="627 376 1406 506">- The table without a primary key lacks a unique identifier for rows. When the network is unstable, you may need to retry the task several times, or data inconsistency may occur.<li data-bbox="627 521 1422 680">- If there are special characters in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database.<li data-bbox="627 696 1422 855">- If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code.<li data-bbox="627 871 1358 972">- After the Oracle table structure is synchronized to the MySQL database, the character set of the table is UTF8MB4.<li data-bbox="627 987 1401 1146">- If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization.<li data-bbox="627 1162 1378 1292">- For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience.<li data-bbox="627 1308 1426 1570">- If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address.<li data-bbox="627 1585 1422 1946">- There are some syntax differences between Oracle and MySQL, so the syntax including but not limited to functions, expressions, and referenced system tables, may not be completely converted during the structure synchronization. Therefore, during the synchronization, the structure may exist in the Oracle database but does not exist in the MySQL database, or the syntax exists in the MySQL database but is not converted. As a result, the structure fails to be synchronized. If this happens, you can manually create a table structure in the destination database.

Type	Restrictions
	<ul style="list-style-type: none"> - In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default. - During full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database. - During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated.
Full synchronization	<ul style="list-style-type: none"> ● When a DRS task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal. ● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases. ● During synchronization, do not perform operations (including but not limited to DDL and DML operations) on the destination database. ● During the synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored. ● During synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails.

Type	Restrictions
Incremental synchronization	<ul style="list-style-type: none">● During synchronization, do not modify or delete the usernames, passwords, permissions, or ports of the source and destination databases.● During synchronization, do not perform operations (including but not limited to DDL and DML operations) on the destination database.● During the synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored.● During synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails.● During synchronization, do not change the char field in the source database or destination database table to varchar, or it is padded with extra spaces due to differences between Oracle and MySQL. In this case, data inconsistency may occur.● During synchronization, some DDL operations are supported.<ul style="list-style-type: none">- Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported.- The object in DDL cannot be the keyword of the destination database, such as index or where. For details about keywords of the destination MySQL database, see MySQL official documentation.- If the destination database version is earlier than 8.0, alter table rename column is not supported.- Incremental DDL operations do not support special characters such as full-width characters.● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization.● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency.● Table names are converted to lowercase letters after the tables are synchronized to the destination database. For example, ABC is converted to abc. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail.● When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be

Type	Restrictions
	synchronized to the destination database. You are advised to add tables during off-peak hours.

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).
- You have read [Suggestions](#) and [Precautions](#).

Procedure

This section uses real-time synchronization from Oracle to RDS for MySQL as an example to describe how to configure a real-time synchronization task.

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Table 4-26 Task and recipient description

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

- Synchronization instance details

Table 4-27 Synchronization instance information

Parameter	Description
Data Flow	Select To the cloud .
Source DB Engine	Select Oracle .
Destination DB Engine	Select MySQL .
Network Type	Public network is used as an example. Available options: VPC , Public network and VPN or Direct Connect

Parameter	Description
Destination DB Instance	The RDS MySQL instance you created.
Synchronization Instance Subnet	<p>Select the subnet where the synchronization instance is located. You can also click View Subnet to go to the network console to view the subnet where the instance resides.</p> <p>By default, the DRS instance and the destination DB instance are in the same subnet. You need to select the subnet where the DRS instance resides and ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.</p>
Synchronization Mode	<ul style="list-style-type: none"> - Full+Incremental This synchronization mode allows you to synchronize data in real time. After a full synchronization initializes the destination database, an incremental synchronization parses logs to ensure data consistency between the source and destination databases. NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible. - Full All database objects and data you selected are synchronized to the destination database at a time. This mode is applicable to scenarios where service interruption is acceptable.

- Enterprise Projects and Tags

Table 4-28 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none"> - This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags. - After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 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**.

Table 4-29 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database. NOTE For a RAC cluster, use a Scan IP address to improve access performance.
Port	The port of the source database. Range: 1 – 65535
Database Service Name	Enter a database service name (Service Name/SID). The client can connect to the Oracle database through the database service name. For details about how to query the database service name, see the prompt on the GUI.
PDB Name	Container database (CDB) and pluggable database (PDB) are new features in Oracle 12c and later versions. This function is optional, but it must be enabled if you want to migrate only PDB tables. Enter the service name, SID, username, and password of the CDB that contains the PDB tables to be migrated.
Database Username	The username for accessing the source database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none">The maximum size of a single certificate file that can be uploaded is 500 KB.If SSL is not enabled, your data may be at risk.

 **NOTE**

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


Table 4-30 Destination database settings

Parameter	Description
DB Instance Name	The RDS for MySQL instance selected when you created the migration task. The instance cannot be changed.
Database Username	The username for accessing the destination database.
Database Password	The database username and password are encrypted and stored in the system and will be cleared after the task is deleted. You can change the password if necessary. To change the password, perform the following operation after the task is created:

Step 4 On the **Set Synchronization Task** page, select the synchronization object type and synchronization objects, and click **Next**.

Table 4-31 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- The flow control mode takes effect only in the full synchronization phase.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Synchronization Object Type	<p>Available options: Table structure, Data, and Constraint</p> <ul style="list-style-type: none">• Data is selected by default.• If Table structure is selected, the destination database cannot contain tables whose names are the same as the source tables to be synchronized.• If Table structure is not selected, the destination database must have tables that match the source tables, and the table structure must be the same as the selected source table structures.
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can synchronize tables or import object files based on your service requirements. To quickly select the desired database objects, you can use the search function.</p> <ul style="list-style-type: none">• If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. <p>NOTE</p> <ul style="list-style-type: none">• To quickly select the desired database objects, you can use the search function.• If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, filter the data to be synchronized and click **Next**. For details, see [Processing Data](#).

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

 **NOTE**


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

Step 7 On the displayed page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

Table 4-32 Task startup settings

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

Step 8 After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper-right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.2.4 From Oracle to PostgreSQL

Supported Source and Destination Databases

Table 4-33 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• On-premises databases• ECS databases	<ul style="list-style-type: none">• RDS for PostgreSQL

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - Start your synchronization 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.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - Data-Level Comparison
To obtain accurate comparison results, **compare data** at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-34 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none"> ● Source database: <ul style="list-style-type: none"> - Full synchronization requires the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, and SELECT for single tables (GRANT SELECT <userName.tbName> to drsUser) - Full+Incremental synchronization: For 12c and later versions, you must have the following permissions: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, SELECT ON SYS.OBJ\$, and SET CONTAINER (GRANT SET CONTAINER TO <userName> CONTAINER=ALL). 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 <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, LOGMINING, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$, 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. For Oracle 11g and earlier versions, the following permissions are required for database synchronization: CREATE SESSION, SELECT ANY DICTIONARY, SELECT for a single table (GRANT SELECT <userName.tbName> to drsUser), EXECUTE_CATALOG_ROLE, SELECT ANY TRANSACTION, SELECT ON SYS.COL\$, and SELECT ON SYS.OBJ\$. - During incremental synchronization, enable PK, UK, 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 destination database user must have the following permissions: INSERT, SELECT, UPDATE, DELETE, CONNECT, CREATE, and REFERENCES.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none">● Only tables, indexes, and (primary key, null, not null) constraints can be synchronized. Views, foreign keys, stored procedures, triggers, functions, events, and virtual columns cannot be synchronized.● In the full synchronization phase, bfile, xml, sdo_geometry, urowid, interval (precision greater than 6 digits), and user-defined types are not supported.● In the incremental synchronization phase, bfile, xml, interval, sdo_geometry, urowid, timestamp (precision greater than 6 digits), and user-defined types are not supported.● During the incremental synchronization, if the source database is a physical standby Oracle database, data of the LOB type cannot be parsed (the data dictionary cannot be generated). If the table to be synchronized contains data of the LOB type, the incremental synchronization will fail.● For the TIMESTAMP WITH TIME ZONE data type, the data cannot be greater than 9999-12-31 23:59:59.999999 after being converted based on the time zone of the destination database.● The default values of the source database can be to_date and sys_guid functions. To use other functions as default values, ensure that the destination database has the same functions. If the destination database does not have the corresponding function, the following results may be displayed:<ul style="list-style-type: none">- The default value may be left empty.- The table fails to be created. As a result, the object comparison result is inconsistent or the task fails.● If the table contains only LOB columns, data inconsistency may occur.● If the empty function of the LOB type is used to write data in the Oracle database, the value queried through JDBC is an empty string. Whether the value is an empty string or NULL after being written to the destination database depends on the processing of the empty string in the destination database.● For a table that does not have a primary key or index, the number of columns of non-large fields must be greater than 3. Otherwise, incremental synchronization may fail because all columns cannot be matched.● Tables whose default values contain expressions of functions cannot be synchronized.● Temporary tables in the source database cannot be synchronized.● When you manually create a table structure, the time type in the destination database must be the same as that in the

Type	Restrictions
	<p>source database. Otherwise, time data may be inconsistent due to time zone conversion.</p>
Source database	<ul style="list-style-type: none"> ● The names of databases and tables cannot contain non-ASCII characters or special characters .><\\,?!" ● An empty source database cannot be synchronized. ● Only the following character sets are supported: ZHS16GBK, AL32UTF8, UTF8, US7ASCII, WE8MSWIN1252, WE8ISO8859P1, WE8ISO8859P2, WE8ISO8859P4, WE8ISO8859P5, WE8ISO8859P7, WE8ISO8859P9, WE8ISO8859P13, WE8ISO8859P15.
Destination database	<ul style="list-style-type: none"> ● The destination DB instance is running properly. ● The destination DB instance must have sufficient storage space. ● Do not use foreign keys for tables during incremental synchronization. Otherwise, the sequence of writing data to different tables may be inconsistent with that in the source database, which may trigger foreign key constraints and cause synchronization failures. ● The destination table can contain more columns than the source table. However, the following failures must be avoided: <ul style="list-style-type: none"> - Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail. - Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.

Type	Restrictions
Precautions	<ul style="list-style-type: none">• Objects that have dependencies must be synchronized at the same time to avoid synchronization failure.• Object names, such as table names, are converted to lowercase letters after being synchronized to the destination database. For example, ABC is converted to abc. In incremental synchronization, the source database cannot contain tables with the same name but different letter cases. Otherwise, the synchronization will fail.• If there are special characters in the Oracle database, the code of the destination Oracle database must be the same as the code of the source Oracle database. Otherwise, garbled characters are displayed in the destination database.• If a task fails because the (self-built) table structure of a table without primary keys is not synchronized and the task cannot be resumed after the table structure is restored, reset the task.• During table structure synchronization, if the destination database has a constraint with the same name as the source database, the table fails to be created.• If you use DRS to synchronize table structures, the tables, constraints, and indexes in the same schema cannot have the same name with different letter cases. For example, table A contains an index named inx1, and table B contains an index name a. Table A and index a have the same name with different letter cases, which will cause table structure synchronization to fail. If multiple schemas are mapped to one schema, the source schemas cannot contain tables, constraints, and indexes with the same name but different letter cases.• If the Oracle character set is WE8MSWIN1252, the CLOB column synchronized to the destination database may contain garbled characters. You can change the character set of the source database to AL32UTF8 before the synchronization.• If the length of a table structure in the Oracle database exceeds 65,535 bytes, the synchronization may fail. The length of a table structure is the total length of all columns. The length of the char or varchar2 type is related to the code.• The size of an Oracle archive log file must be greater than the maximum size of a single data record to prevent incremental data parsing exceptions caused by cross-file (more than two log files) of a single data record.• For an Oracle RAC cluster, use the scan IP address and service name to create a task. The SCAN IP address can provide better fault tolerance, load capability, and synchronization experience.• If the source is an Oracle RAC database and the SCAN IP address is used to configure a DRS task, ensure that the SCAN IP address and DRS node IP address can communicate with all

Type	Restrictions
	<p data-bbox="624 297 1430 461">virtual IP addresses of the source database. Otherwise, the connectivity check fails. If the SCAN IP address is not used, the virtual IP address of a node can be used. In this case, DRS logs are parsed only on the RAC node specified by the virtual IP address.</p> <ul data-bbox="587 477 1430 1984" style="list-style-type: none"><li data-bbox="587 477 1430 607">● If the source is an RAC database, all RAC nodes must be online when incremental synchronization is started for the first time. Otherwise, an error occurs during incremental synchronization.<li data-bbox="587 622 1430 752">● If the source is an RAC database, the number of nodes cannot be increased or decreased during incremental synchronization to avoid incremental synchronization exceptions and ensure strong data consistency.<li data-bbox="587 768 1430 862">● During synchronization, do not delete the username, password, and permissions of the source and destination databases or change the port of the destination database.<li data-bbox="587 878 1430 972">● During the synchronization, do not perform the resetlogs operation on the source Oracle database. Otherwise, data cannot be synchronized and tasks cannot be restored.<li data-bbox="587 987 1430 1081">● During synchronization, the rollback operation of the LOB type is not supported. Otherwise, the synchronization task fails.<li data-bbox="587 1097 1430 1397">● In a full synchronization for the table structure, the length of the char and varchar2 characters in the source database is automatically increased by at least 1.5 times by byte in the destination database (because the length of the destination database is in the unit of byte). The increase multiple depends on the character set of the source and destination databases. For example, if the character set is UTF8, increase the length (byte) by three times by default. If the character set is GBK, increase the length (byte) by two times by default.<li data-bbox="587 1413 1430 1576">● During full synchronization of the partitioned table structure, the table is converted to a non-partitioned table. During incremental synchronization, operations related to the partitioned table in the source database may fail to be executed in the destination database.<li data-bbox="587 1592 1430 1686">● Only normal indexes are synchronized when indexes are synchronized. Primary key constraints are synchronized with the table structure.<li data-bbox="587 1702 1430 1765">● During an incremental synchronization, 0x00 at the end of BLOB and the spaces at the end of CLOB are truncated.<li data-bbox="587 1780 1430 1984">● During incremental synchronization, some DDL operations are supported.<ul data-bbox="624 1854 1430 1984" style="list-style-type: none"><li data-bbox="624 1854 1430 1984">– Table-level synchronization supports alter table add column, alter table drop column, alter table rename column, alter table modify column, and truncate table. The modification of default values is not supported.

Type	Restrictions
	<ul style="list-style-type: none">• The names of mapped table-level objects are case-insensitive. For example, the ABC table mapped to the destination database will be changed to table abc.• When editing the task to add a new table, ensure that transactions of the new table have been committed. Otherwise, transactions that are not committed may fail to be synchronized to the destination database. You are advised to add tables during off-peak hours.• During a full synchronization, DRS writes large amount of data to the destination PostgreSQL database. As a result, the number of PostgreSQL WAL logs increases sharply, and the PostgreSQL disk space may be used up. You can disable the PostgreSQL log backup function before the full synchronization to reduce the number of WAL logs. After the synchronization is complete, enable the function. <p>CAUTION Disabling log backup will affect database disaster recovery. Exercise caution when performing this operation.</p>

4.3 Out of the Cloud

4.3.1 From MySQL to MySQL

Supported Source and Destination Databases

Table 4-35 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• RDS for MySQL	<ul style="list-style-type: none">• On-premises MySQL databases• MySQL databases on an ECS• MySQL databases on other clouds• RDS for MySQL

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).

Suggestions

CAUTION

- When a task is being started or in the full synchronization phase, do not perform DDL operations on the source database. Otherwise, the task may be abnormal.
 - To keep data consistency before and after the synchronization, ensure that no data is written to the destination database during the synchronization.
-
- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
 - Start your synchronization 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. If you have to synchronize data during peak hours, you can select **Yes** for **Flow Control** to adjust the synchronization speed.
 - If network bandwidth is not limited, the query rate of the source database increases by about 50 MB/s during full synchronization, and two to four CPUs are occupied.
 - To ensure data consistency, tables to be synchronized without a primary key may be locked for 3s.
 - The data being synchronized may be locked by other transactions for a long period of time, resulting in read timeout.
 - Due to the inherent characteristics of MySQL, in certain scenarios the performance may be negatively affected. For example, if the CPU resources are insufficient and the storage engine is TokuDB, the read speed on tables may be decreased by 10%.
 - When DRS concurrently reads data from a database, it will use about 6 to 10 sessions. The impact of the connections on services must be considered.
 - If you read a table, especially a large table, during the full migration, the exclusive lock on that table may be blocked.
 - Data-Level Comparison
To obtain accurate comparison results, **compare data** at a specified time point during off-peak hours. If it is needed, select **Start at a specified time** for **Comparison Time**. Due to slight time difference and continuous operations on data, data inconsistency may occur, reducing the reliability and validity of the comparison results.
 - For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail.

Precautions

Before creating a synchronization task, read the following notes:

 NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-36 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none">• The source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT• The destination database user must have the following permissions: SELECT, CREATE, INDEX, DROP, DELETE, INSERT, UPDATE, ALTER, CREATE VIEW, CREATE ROUTINE, and REFERENCES. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.

Type	Restrictions
Synchronization object	<ul style="list-style-type: none">● Only tables, primary key indexes, unique indexes, common indexes, stored procedures, views, and functions can be synchronized.● Table names cannot be mapped for tables on which views, stored procedures, and functions depend.● When table name mapping is used in a synchronization task, foreign key constraints of the table cannot be synchronized.● During database name mapping, if the objects to be synchronized contain stored procedures, views, and functions, these objects cannot be synchronized in the full synchronization phase, resulting in inconsistent objects.● If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters.● Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.● The DDL operation of renaming an unselected table is filtered out during the synchronization. As a result, the task may fail or data may be inconsistent.<ul style="list-style-type: none">- If you rename table A to the name of table B and tables A and B are selected for synchronization, this RENAME statement will not be filtered out.- If you rename table A to the name of table B but table B is not synchronized, this RENAME statement will be filtered out.- You are not advised to perform the rename operation in the many-to-one synchronization scenario. Otherwise, the task may fail or data may be inconsistent.

Type	Restrictions
Source database	<ul style="list-style-type: none">• The source database names cannot contain non-ASCII characters, or the following characters: '<>/'• The source table and view names cannot contain non-ASCII characters, or the following characters: '<>/'• The source database name or mapped name cannot start with ib_logfile or be ib_buffer_pool, ib_doublewrite, ibdata1 or ibtmp1.• During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format.• If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide.• GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.• During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296.

Type	Restrictions
Destination database	<ul style="list-style-type: none">• Data cannot be synchronized from a newer version database to an older version database.• The destination database must have sufficient disk space.• The character set of the destination database must be the same as that of the source database.• The time zone of the destination database must be the same as that of the source database.• If the destination database (excluding MySQL system database) has the same name as the source database, the table structures in the destination database must be consistent with those in the source database.• During a synchronization, a large amount of data is written to the destination database. If the value of the max_allowed_packet parameter of the destination database is too small, data cannot be written. You are advised to set the max_allowed_packet parameter to a value greater than 100 MB.• If the MyISAM tables are included in the synchronization objects, the sql_mode parameter in the destination database cannot contain the no_engine_substitution parameter. Otherwise, the synchronization fails.• The source database names mapped to the destination database cannot contain the following characters: dots (.), angle brackets (<>), backslash (\), and single quotation marks (')

Type	Restrictions
Precautions	<ul style="list-style-type: none">● Objects that have dependencies must be synchronized at the same time to avoid synchronization failure. Common dependencies: tables referenced by views, views referenced by views, views and tables referenced by stored procedures/functions/triggers, and tables referenced by primary and foreign keys● Only MySQL to MySQL synchronization supports many-to-one synchronization. During table-level many-to-one synchronization, tables without primary keys cannot exist in the source database.● If the sources and destinations are RDS instances, database mapping is required.● The source and destination databases cannot contain tables that have the same names but do not have primary keys.● If the source and destination DB instances are RDS MySQL instances, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized.● If the destination MySQL database does not support TLS 1.2 or is a self-built database of an earlier version (earlier than 5.6.46 or between 5.7.0 and 5.7.28), you need to submit an O&M application for testing the SSL connection.● Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key.● The destination table can contain more columns than the source table. However, the following failures must be avoided:<ul style="list-style-type: none">– Assume that extra columns on the destination cannot be null or have default values. If newly inserted data records are synchronized from the source to the destination, the extra columns will become null, which does not meet the requirements of the destination and will cause the task to fail.– Assume that extra columns on the destination must be fixed at a default value and have a unique constraint. If newly inserted data records are synchronized from the source to the destination, the extra columns will contain multiple default values. That does not meet the unique constraint of the destination and will cause the task to fail.● The source database does not support point-in-time recovery (PITR).● The destination database cannot be restored to a point in time when a full synchronization was being performed.● Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.● Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails.

Type	Restrictions
	<ul style="list-style-type: none">• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.• After a task is created, the destination database cannot be set to read-only.• A real-time synchronization task may fail due to the change of the username and password of the source or destination database. If it happens, rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization.• If the source or destination database port is changed during data synchronization, the synchronization task fails. You can rectify the fault as follows:<ul style="list-style-type: none">- If the source database port is wrong, correct the port number on the DRS console and then retry the synchronization task.- If the destination database port is wrong, DRS automatically changes the port to the correct one, and then you need to retry the synchronization task. Generally, do not modify the port number during synchronization.• To ensure data consistency, do not modify the destination database (including but not limited to DDL and DML operations) during synchronization.• Data inconsistency may occur when the MyISAM table is modified during a full synchronization.• Incremental synchronization supports table renaming. Ensure that both the source and destination tables are selected.• You can add additional objects during an incremental synchronization.

Procedure

This section uses synchronization from RDS for MySQL to MySQL as an example to describe how to use DRS to create a real-time synchronization task.

Step 1 On the **Data Synchronization Management** page, click **Create Synchronization Task**.

Step 2 On the **Create Synchronization Instance** page, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Table 4-37 Task and recipient description

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

- Synchronization instance details

Table 4-38 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud . The source database is a database on the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental and Incremental</p> <ul style="list-style-type: none">– 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. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none">– Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

- Enterprise Projects and Tags

Table 4-39 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">– This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.– After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, 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**.

- Source database information

Table 4-40 Source database settings

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

 **NOTE**

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

- Destination database information

Table 4-41 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Port	The port of the destination database. Range: 1 - 65535
Database Username	The username for accessing the destination database.
Database Password	The password for the database username.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none">- The maximum size of a single certificate file that can be uploaded is 500 KB.- If SSL is not enabled, your data may be at risk.


 **NOTE**

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

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Table 4-42 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- The flow control mode takes effect only in the full synchronization phase.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">• 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. <p>In the following scenarios, you can select Ignore or Overwrite. In other scenarios, you are advised to select Report error.</p> <ul style="list-style-type: none">• Data exists in the destination database.• Multiple source databases are synchronized to one destination database.• Data in the destination database is updated manually.

Parameter	Description
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none">• 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.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize data based on service requirements.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none">• If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names.<ul style="list-style-type: none">– If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. <p>NOTE</p> <ul style="list-style-type: none">• To quickly select the desired database objects, you can use the search function.• If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized.• If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed.• The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

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

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

 **NOTE**


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

Step 7 On the displayed page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

Table 4-43 Task startup settings

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

Step 8 After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper-right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.3.2 From MySQL to Kafka

Supported Source and Destination Databases

Table 4-44 Supported databases

Source DB	Destination DB
<ul style="list-style-type: none">• RDS for MySQL	<ul style="list-style-type: none">• Kafka

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).

Suggestions

- The success of database synchronization depends on environment and manual operations. To ensure a smooth synchronization, perform a synchronization trial before you start the synchronization to help you detect and resolve problems in advance.
- It is recommended that you start a task during off-peak hours to minimize the impact of synchronization on your services.

Precautions

Before creating a synchronization task, read the following notes:

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source or destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 4-45 Precautions

Type	Restrictions
Database permissions	<ul style="list-style-type: none">• The source database user must have the following permissions: SELECT, LOCK TABLES, REPLICATION SLAVE, REPLICATION CLIENT, and RELOAD.
Synchronization object	<ul style="list-style-type: none">• The table data can be synchronized.• Tables with storage engine different to MyISAM and InnoDB cannot be synchronized.

Type	Restrictions
Source database	<ul style="list-style-type: none">• During the incremental synchronization, the binlog of the source MySQL database must be enabled and use the row-based format.• If the storage space is sufficient, store the source database binlog for as long as possible. The recommended retention period is three days. If this period is set to 0, the synchronization may fail. If the source database is an RDS for MySQL instance, set the binlog retention period by following the instructions provided in RDS User Guide.• GTID must be enabled for the source database. If GTID is not enabled for the source database, primary/standby switchover is not supported. DRS tasks will be interrupted and cannot be restored during a switchover.• During an incremental synchronization, the server_id value of the MySQL source database must be set. If the source database version is MySQL 5.6 or earlier, the server_id value ranges from 2 to 4294967296. If the source database is MySQL 5.7 or later, the server_id value ranges from 1 to 4294967296.• The database and table names in the source database cannot contain non-ASCII characters, or special characters '<'>/'\
Destination database	<ul style="list-style-type: none">• The destination database is a Kafka database.

Type	Restrictions
Precautions	<ul style="list-style-type: none">• If the data types are incompatible, the synchronization may fail.• If the source DB instance is an RDS MySQL instance, tables encrypted using Transparent Data Encryption (TDE) cannot be synchronized.• A real-time synchronization task may fail due to the change of the username and password of the source database. You need to rectify the information and then retry the synchronization task on the DRS console. Generally, you are advised not to modify the preceding information during synchronization.• If the source database port is changed during data synchronization, the synchronization task fails. If the destination database port is wrong, DRS automatically changes the port to the correct one, and then you need to retry the synchronization task. Generally, do not modify the port number during synchronization.• If a real-time synchronization task fails as the IP address is changed, the system automatically changes the IP address to the correct one. Then, you need to retry the task to continue the synchronization. Therefore, changing the IP address is not recommended.• Cascade operations cannot be performed on tables with foreign keys. If the foreign key index of a table is a common index, the table structure may fail to be created. You are advised to use a unique index.• The source database does not support point-in-time recovery (PITR).• Resumable upload is supported, but data may be repeatedly inserted into a table that does not have a primary key.• Binlogs cannot be forcibly deleted. Otherwise, the synchronization task fails.• The source database does not support the reset master or reset master to command, which may cause DRS task failures or data inconsistency.• Data inconsistency may occur when the MyISAM table is modified during synchronization.• During synchronization of table-level objects, renaming tables is not recommended.• Set the expire_log_day parameter to a proper value to ensure that the binlog does not expire before data transfer resumes. This ensures that services can be recovered after interruption.

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, specify the task name, description, and the synchronization instance details, and click **Create Now**.
- Task information description

Table 4-46 Task and recipient description

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

- Synchronization instance details

Table 4-47 Synchronization instance settings

Parameter	Description
Data Flow	Select Out of the cloud .
Source DB Engine	Select MySQL .
Destination DB Engine	Select Kafka .
Network Type	The Public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Source DB Instance	The RDS for MySQL instance you created.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<ul style="list-style-type: none">– Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.– During synchronization, the source database continues to provide services for external systems with zero downtime.

- Enterprise Projects and Tags

Table 4-48 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">– This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.– After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

NOTE

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

- Step 3** 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**.

Table 4-49 Source database settings

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

 NOTE

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


Table 4-50 Destination database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the destination database.
Security Protocol	Available options: PLAINTEXT , SSL , SASL_PLAINTEXT , and SASL_SSL . For details, see Kafka Authentication .

Step 4 On the **Set Synchronization Task** page, select the synchronization policy, objects, and data format, and click **Next**.

Table 4-51 Synchronization Object

Parameter	Description
Synchronize DML	Select the DML operations to be synchronized. By default, all DML operations are selected.
Topic Synchronization Policy	Topic synchronization policy. You can select A specific topic or Auto-generated topics .
Topic	Select the topic to be synchronized to the destination database. This parameter is available when the topic is set to A specified topic .
Topic Name Format	Topic name format. This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . Only variables database and tablename are supported. The other characters must be constants. Replace \$database\$ with the database name and \$tablename\$ with the table name. For example, if this parameter is set to \$database\$-tablename\$ and the database name is db1 , and the table name is tab1 , the topic name is db1-tab1 . If DDL statements are synchronized, \$tablename\$ is empty and the topic name is db1.
Number of Partitions	This parameter is available when Topic Synchronization Policy is set to Auto-generated topics . The number of partitions of a topic. Each topic can have multiple partitions. More partitions can provide higher throughput but consume more resources. Set the number of partitions based on the actual situation of brokers.

Parameter	Description
Replication Factor	<p>This parameter is available when Topic Synchronization Policy is set to Auto-generated topics.</p> <p>Number of copies of a topic. Each topic can have multiple copies, and the copies are placed on different brokers in a cluster. The number of copies cannot exceed the number of brokers. Otherwise, the topic fails to be created.</p>
Synchronize Topic To	<p>The policy for synchronizing topics to the Kafka partitions.</p> <ul style="list-style-type: none"> • If topics are synchronized to different partitions by hash value of the database and table names, the performance on a single table query can be improved. • If topics are synchronized to partition 0, strong consistency can be obtained but write performance is impacted. If you select Partition 0, only automatically created topics can be synchronized. • If topics are synchronized to different partitions by hash value of the primary key, one table corresponds to one topic.
Data Format in Kafka	<p>Select the data format to be delivered from MySQL to Kafka.</p> <ul style="list-style-type: none"> • Avro refers to binary encoded format. • JSON: JSON message format, which is easy to interpret but takes up more space. • JSON-C: A data format that is compatible with multiple batch and stream computing frameworks. <p>For details, see Kafka Message Format.</p>
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> • If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. <p>NOTE</p> <ul style="list-style-type: none"> • To quickly select the desired database objects, you can use the search function. • If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. • If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. • The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, select the columns to be processed.

- If data processing is not required, click **Next**.
- If you need to process columns, set processing rules by referring to [Processing Data](#).

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


 **NOTE**

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

- Step 7** On the displayed page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

Table 4-52 Task startup settings

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

- Step 8** After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.
- You can view the task status. For more information about task status, see [Task Statuses](#).
 - You can click  in the upper-right corner to view the latest task status.
 - By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.4 Task Management

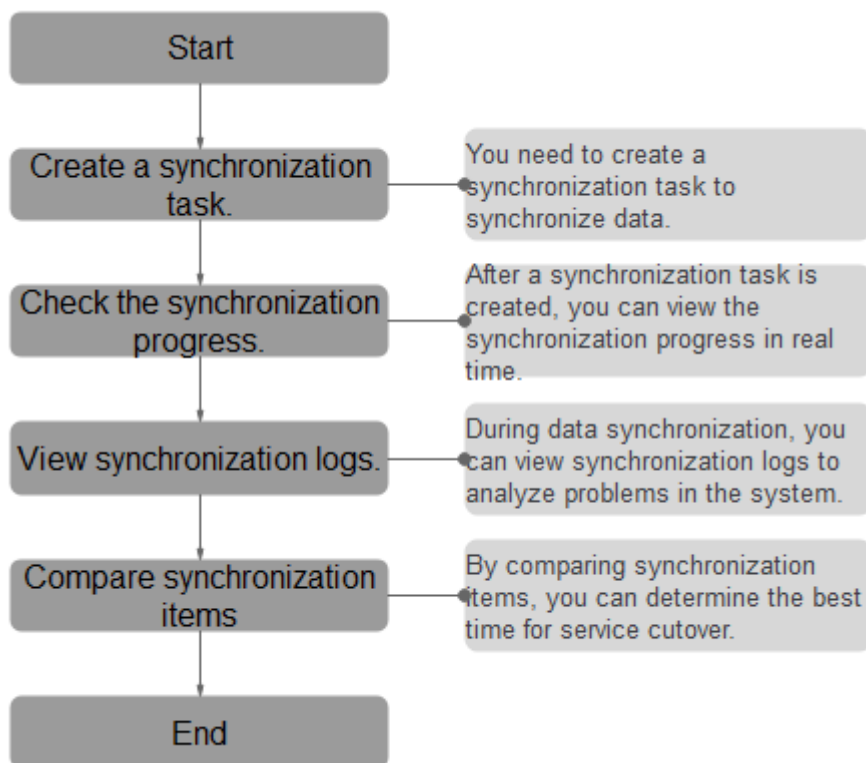
4.4.1 Creating a Synchronization Task

Process

A complete real-time synchronization consists of creating a synchronization task, tracking task progress, analyzing synchronization logs, and comparing data consistency. By comparing multiple items and data, you can synchronize data between different service systems in real time.

A complete real-time synchronization involves the following procedures.

Figure 4-1 Flowchart



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

This section describes how to synchronize data from a MySQL database to an RDS MySQL database. To configure other storage engines, you can refer to the following procedures.

Prerequisites

- You have logged in to the DRS console.
- For details about the DB types and versions supported by real-time synchronization, see [Real-Time Synchronization](#).

Procedure

- Step 1** On the **Data Synchronization Management** page, click **Create Synchronization Task**.
- Step 2** On the **Create Synchronization Instance** page, specify the task name, description, and the synchronization instance details, and click **Create Now**.

- Task information description

Table 4-53 Task and recipient description

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

- Synchronization instance information

Table 4-54 Synchronization instance settings

Parameter	Description
Data Flow	Select To the cloud . The destination database is a database in the current cloud.
Source DB Engine	Select MySQL .
Destination DB Engine	Select MySQL .
Network Type	The public network is used as an example. Available options: Public network, VPC, VPN or Direct Connect
Destination DB Instance	The RDS DB instance you created. NOTE <ul style="list-style-type: none">- The destination DB instance cannot be a read replica.- The source and destination DB instances can be the same DB instance.
Synchronization Instance Subnet	Select the subnet where the synchronization instance is located. 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 ensure that there are available IP addresses. To ensure that the synchronization instance is successfully created, only subnets with DHCP enabled are displayed.

Parameter	Description
Synchronization Mode	<p>Available options: Full+Incremental and Incremental</p> <ul style="list-style-type: none">– 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. <p>NOTE If you select Full+Incremental, data generated during the full synchronization will be continuously synchronized to the destination database, and the source remains accessible.</p> <ul style="list-style-type: none">– Incremental Through log parsing, incremental data generated on the source database is synchronized to the destination database.

- Enterprise Projects and Tags

Table 4-55 Enterprise Projects and Tags

Parameter	Description
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. Select an enterprise project from the drop-down list. The default project is default . For more information about enterprise project, see <i>Enterprise Management User Guide</i> .
Tags	<ul style="list-style-type: none">– This setting is optional. Adding tags helps you better identify and manage your tasks. Each task can have up to 10 tags.– After a task is created, you can view its tag details on the Tags tab. For details, see Tag Management.

 **NOTE**

If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.

Step 3 After the synchronization instance is created, on the **Configure Source and Destination Databases** page, specify source and destination database information. Then, 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**.

- Source database information

Table 4-56 Source database settings

Parameter	Description
IP Address or Domain Name	The IP address or domain name of the source database.
Port	The port of the source database. Range: 1 - 65535
Database Username	The username for accessing the source database.
Database Password	The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created: If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details . In the displayed dialog box, change the password.
SSL Connection	SSL encrypts the connections between the source and destination databases. If SSL is enabled, upload the SSL CA root certificate. NOTE <ul style="list-style-type: none">- The maximum size of a single certificate file that can be uploaded is 500 KB.- If SSL is not enabled, your data may be at risk.

 **NOTE**

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

- Destination database information

Table 4-57 Destination database settings

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

Parameter	Description
Database Password	<p>The password for the database username. You can change the password if necessary. To change the password, perform the following operation after the task is created:</p> <p>If the task is in the Starting, Full synchronization, Incremental synchronization, or Incremental synchronization failed status, in the Connection Information area on the Basic Information tab, click Modify Connection Details. In the displayed dialog box, change the password.</p>

 **NOTE**


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

Step 4 On the **Set Synchronization Task** page, select the conflict policy and synchronization objects, and then click **Next**.

Table 4-58 Synchronization mode and object

Parameter	Description
Flow Control	<p>You can choose whether to control the flow.</p> <ul style="list-style-type: none">• Yes You can customize the maximum migration speed. In addition, you can set the time range based on your service requirements. The traffic rate setting usually includes setting of a rate limiting time period and a traffic rate value. Flow can be controlled all day or during specific time ranges. The default value is All day. A maximum of three time ranges can be set, and they cannot overlap. The flow rate must be set based on the service scenario and cannot exceed 9,999 MB/s.• No The synchronization speed is not limited and the outbound bandwidth of the source database is maximally used, which will increase the read burden on the source database. For example, if the outbound bandwidth of the source database is 100 MB/s and 80% bandwidth is used, the I/O consumption on the source database is 80 MB/s. <p>NOTE</p> <ul style="list-style-type: none">- The flow control mode takes effect only in the full synchronization phase.- You can also change the flow control mode after creating a task. For details, see Modifying the Flow Control Mode.

Parameter	Description
Incremental Conflict Policy	<p>The conflict policy refers to the conflict handling policy during incremental synchronization. By default, conflicts in the full synchronization phase are ignored. Select any of the following conflict policies:</p> <ul style="list-style-type: none">● 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. <p>If data conflicts occur, you can select Ignore, Overwrite, or Report error in any of the following scenarios.</p> <ul style="list-style-type: none">● Data exists in the destination database.● Multiple source databases are synchronized to one destination database.● Data in the destination database is updated manually.
Filter DROP DATABASE	<p>During real-time synchronization, executing DDL operations on the source database may affect the synchronization performance. To reduce the risk of synchronization failure, DRS allows you to filter out DDL operations. Currently, only the delete operations on databases can be filtered by default.</p> <ul style="list-style-type: none">● 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.
Synchronize	<p>Normal indexes and incremental DDLs can be synchronized. You can determine whether to synchronize data based on service requirements.</p>
Start Point	<p>This option is available if you select Incremental in Step 2. The logs of the source database are obtained from the start point during an incremental synchronization.</p> <p>Run show master status to obtain the source database position and set File, Position, and Executed_Gtid_Set as prompted.</p>

Parameter	Description
Synchronization Object	<p>The left pane displays the source database objects, and the right pane displays the selected objects. You can select Tables or Databases for Synchronization Object as required.</p> <ul style="list-style-type: none"> If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. For details, see Mapping Object Names. <ul style="list-style-type: none"> If the database table name contains characters other than letters, digits, and underscores (_), or the mapped database table name contains hyphens (-) and number signs (#), the name length cannot exceed 42 characters. <p>NOTE</p> <ul style="list-style-type: none"> To quickly select the desired database objects, you can use the search function. If there are changes made to the source databases or objects, click  in the upper right corner to update the objects to be synchronized. If the object name contains spaces, the spaces before and after the object name are not displayed. If there are multiple spaces between the object name and the object name, only one space is displayed. The name of the selected synchronization object cannot contain spaces.

Step 5 On the **Process Data** page, set the filtering rules for data processing.

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

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

 **NOTE**


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

Step 7 On the displayed page, specify **Start Time**, confirm that the configured information is correct, and click **Submit** to submit the task.

Table 4-59 Task startup settings

Parameter	Description
Started Time	<p>Set Start Time to Start upon task creation or Start at a specified time based on site requirements.</p> <p>NOTE 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.</p>

Step 8 After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.

- You can view the task status. For more information about task status, see [Task Statuses](#).
- You can click  in the upper-right corner to view the latest task status.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.

----End

4.4.2 Querying the Synchronization Progress

This section describes how to check the synchronization progress.

- During a full synchronization, DRS displays the progress overview. You can view the structure, data, and index migration progress. When the progress reaches 100%, the synchronization is complete. The synchronization of data and indexes is relatively slow.
- During an incremental synchronization, DRS displays the incremental synchronization delay. You can determine the synchronization status between the source and destination databases based on the delay. If the delay is 0, the source and destination databases are instantaneously consistent, and no new transaction needs to be synchronized.

Prerequisites

You have logged in to the DRS console.

Procedure

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

Step 2 On the displayed page, click **Synchronization Progress** to view table synchronization progress.

- When a full synchronization is complete, the progress reaches 100%.
- After the full synchronization is complete, the incremental synchronization starts. You can view the incremental synchronization delay on the **Synchronization Progress** tab.

- You can also view the incremental synchronization delay on the **Data Synchronization Management** page. When the incremental synchronization delay exceeds the preset or default threshold, the value of the incremental synchronization delay is displayed in red in the task list.
- When the delay is 0s, the data in the source and destination databases is synchronized in real time.

NOTE

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

Transactions are synchronized as follows:

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

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

CAUTION

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

----End

4.4.3 Viewing Synchronization Logs

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

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click **Synchronization Logs** to view the logs generated during the synchronization.

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

----End

4.4.4 Comparing Synchronization Items

Scenarios

This section describes how to compare synchronization items to check if there are any differences between source and destination databases. To minimize the impact on services and shorten the service interruption duration, the following comparison methods are provided:

- Object-level comparison: compares objects such as databases, indexes, tables, views, stored procedures, functions, and table sorting rules.
- Data-level comparison is classified into row comparison and value comparison.
 - Row comparison: It helps you compare the number of rows in the tables to be synchronized. This comparison method is recommended because it is fast.
 - Value comparison: It helps you check whether data in the synchronized table is consistent. The comparison process is relatively slow.

When you check data consistency, compare the number of rows first. If the number of rows are inconsistent, you can then compare the data in the table to determine the inconsistent data.

Constraints

- A comparison task can be created only when the task is in the incremental phase. When a full task is complete, DRS automatically creates object-level and row comparison tasks.
- If DDL operations were performed on the source database, you need to compare the objects again to ensure the accuracy of the comparison results.
- If data in the destination database is modified separately, the comparison results may be inconsistent.
- Currently, only tables with primary keys support value comparison. For tables that do not support value comparison, you can compare rows. Therefore, you can compare data by row or value based on scenarios.
- To prevent resources from being occupied for a long time, DRS limits the row comparison duration. If the row comparison duration exceeds the threshold, the row comparison task stops automatically. If the source database is a relational database, the row comparison duration is 60 minutes. If the source database is a non-relational database, for example, MongoDB, the row comparison duration is 30 minutes.
- In the many-to-one row comparison scenario, the number of rows in the table in the source database is compared with that in the aggregation table mapped to the destination database.
- If the source is a PostgreSQL database, the index and constraint names will be changed during table mapping. As a result, the index and constraint names are inconsistent.

Prerequisites

- You have logged in to the DRS console.

- A synchronization task has been started.

Creating a Comparison Task

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

Step 2 Click the **Synchronization Comparison** tab.

Step 3 Compare synchronization items.

- On the **Object-Level Comparison** tab, check whether the comparison results of the source and destination databases are consistent. Locate a comparison item you want to view and click **View Details** in the **Operation** column.
- On the **Data-Level Comparison** tab, click **Create Comparison Task**. In the displayed dialog box, specify **Comparison Type**, **Comparison Time**, and **Object**. Then, click **OK**.

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

- Row comparison: checks whether the source table has the same number of rows as the destination table.

NOTE

- After a task enters the incremental comparison phase, you can create a row comparison task.

- Value comparison: checks whether the source table has the same data as the destination table.

NOTE

- After a task enters the incremental synchronization phase, you can create a value comparison task. After the full synchronization is complete, data in the source database cannot be changed. Otherwise, the comparison result will be inconsistent.

Value comparison only applies to tables with single-column primary key or unique index. You can use row comparison for tables that do not support value comparison. Therefore, you can compare data by row or value based on scenarios.

– **Comparison Policy:** DRS supports one-to-one and many-to-one comparison policies.

- **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.


NOTE

If you select **Row Comparison** for **Comparison Type**, the **Comparison Policy** option becomes available.

– **Comparison Time:** You can select **Start upon task creation** or **Start at a specified time**. There is a slight difference in time between the source and destination databases during synchronization. Data inconsistency

may occur. You are advised to compare migration items during off-peak hours for more accurate results.

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

Step 4 After the comparison creation task is submitted, the **Data-Level Comparison** tab is displayed. Click  to refresh the list and view the comparison result of the specified comparison type.

Value comparison only applies to tables with single-column primary key or unique index. You can use row comparison for tables that do not support value comparison. Therefore, you can compare data by row or value based on scenarios.

If you want to view the row or value comparison details, click **View Results**.

If you want to download the row comparison or value comparison result, locate a specified comparison type and click **Export Report** in the **Operation** column.

 **NOTE**

You can also view comparison details of canceled comparison tasks.

----End

4.4.5 Managing Objects

4.4.5.1 Editing Synchronization Objects

This section describes how to change synchronization objects in a synchronization task. After a data synchronization task is created, you can change synchronization objects by adding or deleting databases and tables to be synchronized during the incremental synchronization.

Prerequisites

You have logged in to the DRS console.

Method 1

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

Step 2 On the **Set Synchronization Task** page, change the objects to be synchronized and click **Next**.

Step 3 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**.

 **NOTE**

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

Step 4 In the synchronization task list on the **Data Synchronization Management** page, the current task status is **Incremental synchronization**, and a subtask in the **Modifying task** status is generated. After the subtask change is complete, incremental synchronization is performed for the edited synchronization objects.

----End

Method 2

Step 1 On the **Data Synchronization Management** page, click the target synchronization task.

Step 2 On the displayed page, click the **Synchronization Mapping** tab and click **Edit** to the right of the synchronization object.

Step 3 Perform [Step 2](#) to [Step 4](#) from method 1.

----End

4.4.5.2 Mapping Object Names

Data synchronization allows you to synchronize objects (including databases, schemas and tables) in a sources database to the corresponding objects in a destination database. If the synchronization objects in source and destination databases have different names, you can map the source object name to the destination one. The object types that can be mapped include database, schema, and table.

Object name mapping can be used only in the following scenarios:

- For the first time you select synchronization objects for a data synchronization task.
- For the first time you add or delete the synchronization object which is not in a mapping relationship.

NOTE

- If the destination DB is a type of PostgreSQL (for example, GaussDB, or PostgreSQL), the destination schema name cannot start with **pg_**. Otherwise, the migration fails.

This section describes how to map objects when configuring a data synchronization task. For details about the mapping relationship, see [Viewing Synchronization Mapping Information](#).

Mapping Databases

During real-time synchronization, if the names of source databases to be synchronized are different from those in the destination, you can map the source database names to the destination ones.

Step 1 On the **Set Synchronization Task** page, select the database that needs to be mapped from the synchronization objects on the right area and click **Edit**.

Step 2 Changing a database name

In the displayed dialog box, enter a new database name. The new database name is the name of the database saved in the destination DB instance.

Step 3 Check the result.

After the database name is changed, the database name before modification and the new database name are displayed. The database mapping is complete.

----End

Mapping Schemas

A schema is a collection of database objects, including: tables, views, stored procedures, and indexes.

During real-time synchronization, if the names of source schemas to be synchronized are different from those in the destination, you can map the source schema names to the destination ones.

Step 1 On the **Set Synchronization Task** page, select the schema that needs to be mapped from the synchronization objects on the right area and click **Edit**.

Step 2 Edit the schema name.

In the displayed dialog box, enter a new schema name which is the name to be saved in the destination database.

Step 3 Check the result.

After the schema name is changed, the schema name before modification and the new schema name are displayed. The schema mapping is complete.

----End

Mapping Tables

During real-time synchronization, if the names of source tables to be synchronized are different from those in the destination, you can map the source table names to the destination ones.

Step 1 On the **Set Synchronization Task** page, select the table that needs to be mapped from the synchronization objects on the right area and click **Edit**.

Step 2 Change a table name.

In the displayed dialog box, enter a new table name. The new table name is the name of the table saved in the destination database.

Step 3 Check the result.

After the table name is changed, the table name before modification and the new table name are displayed. The table mapping is complete.

----End

4.4.5.3 Viewing Synchronization Mapping Information

During real-time synchronization, the objects that can be mapped to the destination include databases, schemas, tables, and columns (in data processing). After a mapping relationship between objects is established, you can view details about the mapping.

Prerequisites

You have logged in to the DRS console.

Procedure

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task name in the **Task Name/ID** column.
- Step 2** On the displayed page, click the **Synchronization Mapping** tab to view the mapping details.

NOTE

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

- Step 3** In the upper right corner, filter and search for the mapping relationships by object or column.

----End

4.4.5.4 Processing Data

DRS processes synchronized objects and allows you to add rules for selected objects. The processing rules supported by each data flow type are different.

Adding Additional Columns

- Step 1** On the **Process Data** page of the real-time synchronization task, click **Additional Columns**, locate the table to be processed, and click **Add** in the **Operation** column.
- Step 2** In the displayed **Add** dialog box, specify the column name, operation type, and field type. Click **OK**.

NOTE

- In many-to-one mapping scenarios, additional columns for data processing are required to avoid data conflicts.
- The following operation types are supported:
 - **Default:** Use the default value to fill in the new column.
 - Use the `create_time` column and `update_time` column as an example to fill the new column with the data creation time and data update time.
 - **Expression:** indicates that you need to fill the new column with an expression.
 - If you fill in the new column in `serverName@database@table` format, you need to enter a server name and then the database name and table name will be automatically filled in.
 - **Value:** Select a value, for example, synchronization time.
- You can apply the additional column information of the first editable table to all editable tables in batches.
- During MySQL to GaussDB(for MySQL) primary/standby synchronization, if the number of columns in a single table exceeds 500, the number of additional columns added to the table may exceed the upper limit. As a result, the task fails.

Step 3 Click **Next**.

----End

Filtering Data

After a data filtering rule is added, update the source database to ensure data consistency. For example:

- The filter criteria are met after the update. You need to continue the synchronization and perform the same update operation on the destination database. If no data is matched, the operation will be ignored, causing data inconsistency.
- The filter criteria are not met after the update. You need to continue the synchronization and perform the same update operation on the destination database.

Step 1 On the **Processing Data** page, set **Processing Type** to **Data filtering**.

Step 2 In the **Object** area, select the table to be processed.

Step 3 In the **Filtering Criteria** area, enter the filter criteria (only the part after WHERE in the SQL statement, for example, id=1), and click **Verify**.

NOTE

- Each table has only one verification rule.
- Up to 500 tables can be filtered at a time.
- The filter expression cannot use the package, function, variable, or constant of a specific DB engine. It must comply with the general SQL standard. Enter the part following WHERE in the SQL statement (excluding WHERE and semicolons), for example, sid > 3 and sname like "G %". A maximum of 512 characters are allowed.
- In SQL statements for setting filter criteria, keywords must be enclosed in backquotes, and the value of **datetime** (including date and time) must be enclosed in single quotation marks, for example, `update` > '2022-07-13 00:00:00' and age >10.
- Filter criteria cannot be configured for large objects, such as CLOB, BLOB, and BYTEA.
- You are not advised to set filter criteria for fields of approximate numeric types, such as FLOAT, DECIMAL, and DOUBLE.
- Do not use fields containing special characters as a filter condition.
- You are not advised to use non-idempotent expressions or functions as data processing conditions, such as SYSTIMESTAMP and SYSDATE, because the returned result may be different each time the function is called.

Step 4 After the verification is successful, click **Generate Processing Rule**. The rule is displayed.

Step 5 Click **Next**.

----End

Advanced Settings for Data Filtering

If you need to query an association table, you can use the advanced settings of data processing.

Step 1 On the **Process Data** page of the real-time synchronization task, set **Processing Type** to **Data filtering**.

Step 2 In the **Object** area, select the table to be processed.

Step 3 In the **Filtering Criteria** area, specify the filtering criteria, for example, id1 in (select id from db1.tab1 where id >=3 and id <10), and click **Verify**.

 **NOTE**

- Each table has only one verification rule.
- Up to 500 tables can be filtered at a time.
- The filter expression cannot use the package, function, variable, or constant of a specific DB engine. It must comply with the general SQL standard. Enter the part following WHERE in the SQL statement (excluding WHERE and semicolons), for example, sid > 3 and sname like "G %". A maximum of 512 characters are allowed.
- Filter criteria cannot be configured for large objects, such as CLOB, BLOB, and BYTEA.
- You are not advised to set filter criteria for fields of approximate numeric types, such as FLOAT, DECIMAL, and DOUBLE.
- Do not use fields containing special characters as a filter condition.
- You are not advised to use non-idempotent expressions or functions as data processing conditions, such as SYSTIMESTAMP and SYSDATE, because the returned result may be different each time the function is called.

Step 4 After the verification is successful, click **Generate Processing Rule**. The rule is displayed.

Step 5 In the **Advanced Settings** area, specify the configuration condition and rule for the association table to help you filter data.

1. In the **Configuration Condition** area, enter the association table information entered in [Step 3](#).

Database Name, Table Name, Column Name, Primary Key, Index, and Filter Criteria are mandatory. If the table does not have an index, enter its primary key.

Filter Criteria is the filter condition of the association table information entered in [Step 3](#).

2. Then, click **Verify**.
3. After the verification is successful, click **Generate Configuration Rule**. The rule is displayed in the **Configuration Rule** area.

To filter data in multiple association tables, repeat [Step 5](#).

 **NOTE**

Configuration rules can be deleted.

Step 6 Click **Next**.

----End

Processing Columns

Step 1 On the **Process Data** page of the real-time synchronization task, select **Processing Columns**.

Step 2 In the **Object** area, select the objects to be processed.

Step 3 Click **Edit** to the right of the selected object.

Step 4 In the **Edit Column** dialog box, select the columns to be mapped and enter new column names.

 **NOTE**

- You can query or filter columns or create new column names.
- After the column name is edited, the column name of the destination database is changed to the new name.
- The new column name cannot be the same as the original column name or an existing column name.
- The column name in the synchronized table cannot be modified.
- Only the selected columns can be synchronized.
- MySQL to MySQL synchronizations do not support column mapping based on the partitioning column of a partitioned table.


Step 5 Click **Confirm**.

Step 6 Click **Next**.

----End

Viewing Data Filtering Results


Step 1 In the task list, click the task to be processed.

Step 2 Click the **Process Data** tab to view data filtering records. Click  in the upper right corner to refresh the record list.

----End

View Column Processing

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

Step 2 In the navigation pane on the left, choose **Synchronization Mapping**. In the upper right corner, and select **Columns** to view column mapping records. Click  in the upper right corner to refresh the record list.

----End

4.4.6 Task Life Cycle

4.4.6.1 Viewing Task Details

View the information about the synchronization tasks and synchronization instances. This section describes how to view details about a synchronization task you have created.

Prerequisites

You have logged in to the DRS console.

Procedure

NOTE

In the task list, only tasks created by the current login user are displayed. Tasks created by different users of the same tenant are not displayed.

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

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

You can view information about the task, synchronization instance, and synchronization.

----End

4.4.6.2 Modifying Task Information

After a synchronization task is created, you can modify task information to identify different tasks.

The following task information can be edited:

- Task name
- Description
- Task start time

Prerequisites

You have logged in to the DRS console.

Procedure

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

Step 2 On the **Basic Information** tab, locate the information to be modified in the **Task Information** area.




- You can click  to modify the task name and description.
 - To submit the change, click .
 - To cancel the change, click .

Table 4-60 Task information

Task Information	Description
Task Name	The task name must start with a letter and consist of 4 to 50 characters. It can contain only letters, digits, hyphens (-), and underscores (_).

Task Information	Description
Description	The description consists of a maximum of 256 characters and cannot contain the following special characters: !<>&'\"

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

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

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

----End

4.4.6.3 Modifying Connection Information

During the synchronization, you may change the password of the source or destination database. As a result, the data synchronization, data comparison, task resuming, resetting, object editing, and stopping may fail. In this case, you need to change the password on the DRS console and resume the task.

You can modify the following synchronization information:

- Source database password
- Destination database password

NOTE

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

Prerequisites

You have logged in to the DRS console.

Procedure

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

Step 2 On the **Basic Information** tab, click **Modify Connection Details** in the **Connection Information** area.

Step 3 In the displayed dialog box, change the passwords of the source and destination databases and click **OK**.

----End

4.4.6.4 Modifying the Flow Control Mode

You can choose whether to control the flow. DRS allows you to change the flow control mode after a task is created. Currently, only the real-time synchronization scenarios listed in [Real-time Synchronization Scenarios That Support Flow Control](#) support this function.

 **NOTE**

- The flow control mode takes effect only in the full synchronization phase.
- After the traffic rate is modified in the incremental migration phase, the modification takes effect when the task enters the full migration phase again. For example, if the traffic rate is modified and a synchronization object is added to the task, the modification takes effect in the full synchronization phase of the task.

Prerequisites

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

Method 1

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

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

----End

Method 2

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

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

----End

Real-time Synchronization Scenarios That Support Flow Control

To the cloud

- MySQL->MySQL
- PostgreSQL->PostgreSQL

From the cloud

- MySQL->MySQL

4.4.6.5 Editing a Synchronization Task

For a synchronization task that has been created but not started, DRS allows you to edit the configuration information of the task, including the source and destination database details. For synchronization tasks in the following statuses, you can edit and submit the tasks again.

- Creating
- Configuration

 **NOTE**

For an incremental synchronization task, DRS allows you to modify synchronization objects. For details, see [Editing Synchronization Objects](#).

Prerequisites

You have logged in to the DRS console.

Method 1

- Step 1** In the task list on the **Data Synchronization Management** page, locate the target task and click **Edit** in the **Operation** column.
- Step 2** On the **Configure Source and Destination Databases** page, enter information about the source and destination databases and click **Next**.
- Step 3** On the **Set Synchronization Task** page, select synchronization objects and click **Next**.
- Step 4** On the **Check Task** page, check the synchronization task.
- Step 5** On the **Confirm Task** page, specify **Start Time**, confirm that the configured information is correct, and click **Next**.
- Step 6** After the task is submitted, you can view and manage it on the **Data Synchronization Management** page.

----End

Method 2

- Step 1** On the **Data Synchronization Management** page, click the target synchronization task.
- Step 2** On the displayed page, click **edit this task** to go to the **Configure Source and Destination Databases** page.
- Step 3** Perform [Step 2](#) to [Step 6](#) from method 1.

----End

4.4.6.6 Resuming a Synchronization Task

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

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

- Synchronization failed
- Paused

NOTE

- If the synchronization task fails due to non-network problems, the system will automatically resume the task three times by default. If the failure persists, you can resume the task manually.
- If the synchronization fails due to network problems, the system will automatically resume the task until the synchronization is restored.

Prerequisites

You have logged in to the DRS console.

Method 1

In the task list on the **Data Synchronization Management** page, locate the target task and click **Resume** in the **Operation** column.

Method 2

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

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

----End

Resume Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be resumed.

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

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

----End

4.4.6.7 Pausing a Synchronization Task

DRS allows you to pause real-time synchronization tasks. For details about the synchronization scenarios where synchronization tasks can be paused, see [Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Paused](#).

Prerequisites

- You have logged in to the DRS console.

Pausing a Task

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

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

NOTE

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

----End

Pausing Tasks

Step 1 On the **Data Synchronization Management** page, select the tasks to be paused.

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

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

----End

Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Paused

The following tasks can be paused during incremental synchronization:

- To the cloud
 - MySQL->MySQL
 - MySQL->GaussDB(DWS)
 - PostgreSQL->PostgreSQL
 - Oracle->GaussDB(DWS)
- From the cloud
 - MySQL->MySQL
 - MySQL->Kafka

In addition, the following tasks can be paused during full synchronization:

- MySQL->MySQL

4.4.6.8 Resetting a Synchronization Task

During real-time synchronization, you can reset the synchronization tasks in one of the following statuses so that you do not need to configure the tasks again.

- Paused
- Failed

For details about the synchronization scenarios where synchronization tasks can be reset, see [Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Reset](#).

NOTE

- For a many-to-one synchronization task, only the parent task can be reset.
- For a MySQL many-to-one synchronization task, only the subtask can be reset.
- Resetting a task does not clear the destination database. You can determine whether to clear the destination database based on your service requirements. After the task is reset, a full synchronization is performed again. You do not need to configure the task again.

Prerequisites

You have logged in to the DRS console.

Method 1

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

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

 **NOTE**

If a many-to-one synchronization task fails to be reset, click the name of the failed subtask in the failure details to view the failure cause of the task.

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

----End

Method 2

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

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

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

----End

Real-time Synchronization Scenarios Where Synchronization Tasks Can Be Reset

To the cloud

- MySQL->MySQL
- MySQL->GaussDB(DWS)
- PostgreSQL->PostgreSQL
- Oracle -> GaussDB(DWS) (Only tasks in the failed state can be reset.)

From the cloud

- MySQL->MySQL
- MySQL -> Kafka (Only tasks in the failed state can be reset.)

4.4.6.9 Stopping a Synchronization Task

After the source database and services are migrated to the destination database, you can stop the synchronization task. To prevent data from being overwritten after the source database and services are migrated to the destination database, stop a synchronization task to achieve this goal.

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

- Creating
- Configuration
- Pending start

- Full synchronization
- Full synchronization failed
- Incremental synchronization
- Incremental synchronization failed
- Paused
- Fault rectification

NOTICE

- You are advised to stop the task before performing other operations, such as disconnecting the network between the source database and the synchronization instance. Otherwise, an alarm indicating that the source database cannot be connected will be generated.
 - For a task in the **Configuration** state, it cannot be stopped if it fails to be configured.
 - For a task in the **Fault rectification** state, it cannot be stopped if the fault is being rectified.
 - After a task is stopped, it cannot be retried.
-

Procedure

Step 1 In the task list on the **Data Synchronization Management** page, locate the target task and click **Stop**.

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

 **NOTE**

- If the task status is abnormal (for example, the task fails or the network is abnormal), DRS will select **Forcibly stop task** to preferentially stop the task to reduce the waiting time.
- Forcibly stopping a task will release DRS resources. Check whether the synchronization is affected.
- To stop the task properly, restore the DRS task first. After the task status becomes normal, click **Stop**.

----End

4.4.6.10 Deleting a Synchronization Task

This section describes how to delete a synchronization task that has been completed or has failed. Deleted tasks will no longer be displayed in the task list. Exercise caution when performing this operation.

Prerequisites

You have logged in to the DRS console.

Deleting a Task

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

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

----End

4.4.6.11 Task Statuses

Synchronization statuses indicate different synchronization phases.

Table 4-61 lists synchronization task statuses and descriptions.

Table 4-61 Task status description

Status	Description
Creating	A synchronization instance is being created.
Task creation failed	Failed to create a real-time synchronization instance.
Configuration	The synchronization instance is successfully created, but the synchronization task is not started. You can continue to configure the task.
Pending start	The scheduled synchronization task has been delivered to the synchronization instance, waiting for the synchronization instance to start the synchronization task.
Starting	The task is being started.
Start failed	A real-time synchronization task fails to be started.
Full synchronization	A full synchronization task is being performed.
Full synchronization failed	A full synchronization task fails.
Incremental synchronization	An incremental synchronization task is being performed.
Incremental synchronization failed	An incremental synchronization task fails.
Modifying task	The synchronization object is being modified.
Modifying task failed	The synchronization object fails to be modified.

Status	Description
Fault rectification	A synchronization instance is faulty and the system automatically restores the synchronization task.
Paused	The real-time synchronization task has been paused.
Task stopping	The synchronization instance and resources used for executing the synchronization task are being released.
Completing	A synchronization instance and resources are being released.
Stopping task failed	The synchronization instance and resources fail to be released.
Completed	The task is completed and the synchronization instance is released.

 **NOTE**

- If a task fails to be created, DRS retains the task for three days by default. After three days, the task automatically ends.
- By default, DRS retains a task in the **Configuration** state for three days. After three days, DRS automatically deletes background resources, and the task status remains unchanged. When you restart the task configuration, DRS applies for resources again.
- Deleted synchronization tasks are not displayed in the status list.

4.5 Tag Management

Scenarios

Tag Management Service (TMS) enables you to use tags on the management console to manage resources. TMS works with other cloud services to manage tags. TMS manages tags globally, and other cloud services manage their own tags. If you have to manage a large number of tasks, you can use different tags to identify and search for tasks.

- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- Each DB instance can have up to 10 tags.

Adding a Tag

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

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

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

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

----End

Editing a Tag

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

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

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

----End

Delete a Tag

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

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

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

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

----End

4.6 Interconnecting with Cloud Eye

4.6.1 Supported Metrics

Description

This section describes metrics reported by the Data Replication Service (DRS) to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated for DRS.

Namespace

SYS.DRS

DB Instance Monitoring Metrics

[Table 4-62](#) lists the DRS performance metrics.

Table 4-62 DRS metrics

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

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

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

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

Dimensions

Key	Value
instance_id	DRS instance ID

4.6.2 Configuring Alarm Rules

Scenarios

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

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

Procedure

- Step 1** Log in to the management console.
 - Step 2** Under **Management & Deployment**, click **Cloud Eye**.
 - Step 3** In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
 - Step 4** On the displayed **Alarm Rules** page, click **Create Alarm Rule**.
- End

4.6.3 Viewing Monitoring Metrics

Scenarios

Cloud Eye monitors the running statuses of replication, synchronization, and DR instances. You can obtain the monitoring metrics on the management console. Monitored data requires a period of time for transmission and display. The status of the monitored object displayed on the Cloud Eye page is the status obtained 5 to 10 minutes before. You can view the monitored data of a newly created DB instance 5 to 10 minutes later.


Prerequisites

An instance is running properly when in the following statuses:

- Real-time migration: Full migration and Incremental migration
- Real-time synchronization: Full synchronization and Incremental synchronization
- Real-time disaster recovery: Disaster recovery in progress

Viewing Metrics

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and project.

Step 3 Choose **Database > Data Replication Service**. The **Data Replication Service** page is displayed.

Step 4 Take real-time migration as an example. On the **Online Migration Management** page, click the target migration task name in the **Task Name/ID** column.

Step 5 On the displayed page, click **View Metric** in the upper right corner of the page to go to the Cloud Eye console.

By default, the monitoring information about the DRS instance is displayed on this page.

Step 6 View monitoring metrics of the instance.

- On the Cloud Eye console, click the target DB instance name and click **Select Metric** in the upper right corner. In the displayed dialog box, you can select the metrics to be displayed and sort them by dragging them at desired locations.
- You can sort graphs by dragging them based on service requirements.
- Cloud Eye can monitor performance metrics from the last 1 hour, 3 hours, 12 hours, 1 day, 7 days, and 6 months.

----End

4.7 Operation Reference in Synchronization Scenarios

4.7.1 Kafka Message Format

Data synchronized to the Kafka cluster is stored in Avro, JSON, and JSON-C formats.

Avro Format

For details about the schema definition in Avro format, see [record.rar](#). After data is synchronized to the Kafka cluster, parse data based on the definition of the Avro schema. For details about the data parsing, see [drs-avro-read.rar](#).

JSON

For details about the JSON format from MySQL to Kafka, see [Table 4-63](#).

Table 4-63 Parameters for synchronizing from MySQL to Kafka

Parameter	Description
mysqlType	Field name and type in the source table.

Parameter	Description
id	Sequence number of an event operation defined in DRS. The value increases monotonically.
es	The time when the record is generated in the source database. The value is a 13-digit Unix timestamp in milliseconds.
ts	The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds.
database	Database name
table	Table name.
type	Operation type, such as DELETE, UPDATE, INSERT, and DDL.
isDdl	Whether the operation is a DDL operation.
sql	A DDL-defined SQL statement. The value is "".
sqlType	JDBC type of the fields in the source table.
data	The latest data, which is a JSON array. If the value of type is INSERT , the latest data is inserted. If the value of type is UPDATE , the latest data is updated.
old	Old data. If the value of type is UPDATE , the data is old. If the value of type is DELETE , the data is deleted.
pkNames	Primary key name

```
{
  "mysqlType":{
    "c11":"binary",
    "c10":"varchar",
    "c13":"text",
    "c12":"varbinary",
    "c14":"blob",
    "c1":"varchar",
    "c2":"varbinary",
    "c3":"int",
    "c4":"datetime",
    "c5":"timestamp",
    "c6":"char",
    "c7":"float",
    "c8":"double",
    "c9":"decimal",
    "id":"int"
  },
  "id":27677,
  "es":1624614713000,
  "ts":1625058726990,
  "database":"test01",
  "table":"test ",
  "type":"UPDATE",
  "isDdl":false,
  "sql": "",
  "sqlType":{
```

```
"c11":-2,  
"c10":12,  
"c13":-1,  
"c12":-3,  
"c14":2004,  
"c1":12,  
"c2":-3,  
"c3":4,  
"c4":94,  
"c5":93,  
"c6":1,  
"c7":6,  
"c8":8,  
"c9":3,  
"id":4  
},  
"data":[  
  {  
    "c11": "",  
"c10": "cloud",  
  
"c13": "asfiahfiaf939-0239uoituqorjqirfoidfqrniowejoiwjqroqwjrowqjoioiqgoiegnkjgoi23roiugou  
ofdug9u90weurtg103",  
"c12": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111,  
103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48,  
57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 101, 117, 114, 103, 57,  
101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114, 48, 52, 117, 48, 57, 53, 116, 117, 51, 48,  
57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",  
"c14": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111,  
103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48,  
57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 55, 57, 56, 52, 54, 53,  
52, 54, 54, 54, 49, 52, 54, 53, 33, 64, 35, 36, 37, 94, 42, 40, 41, 95, 41, 43, 95, 43, 124, 125, 34,  
63, 62, 58, 58, 101, 117, 114, 103, 57, 101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114,  
48, 52, 117, 48, 57, 53, 116, 117, 51, 48, 57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",  
"c1": "cf3f70a7-7565-44b0-ae3c-83bec549ea8e:104",  
"c2": "",  
"c3": "103",  
"c4": "2021-06-25 17:51:53",  
"c5": "1624614713.201",  
"c6": "!@#%$90weurtg103",  
"c7": "10357.0",  
"c8": "1.2510357E7",  
"c9": "9874510357",  
"id": "104"  
  }  
],  
"old":[  
  {  
    "c11": "",  
"c10": "cloud",  
  
"c13": "asfiahfiaf939-0239",  
"c12": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111,  
103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48,  
57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 101, 117, 114, 103, 57,  
101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114, 48, 52, 117, 48, 57, 53, 116, 117, 51, 48,  
57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",  
"c14": "[106, 103, 111, 106, 103, 111, 105, 100, 115, 106, 103, 111, 106, 111, 115, 111,  
103, 57, 51, 52, 48, 57, 52, 51, 48, 57, 116, 106, 104, 114, 103, 106, 101, 119, 57, 116, 117, 48,  
57, 51, 52, 48, 116, 101, 114, 111, 101, 106, 103, 57, 56, 51, 48, 52, 105, 55, 57, 56, 52, 54, 53,  
52, 54, 54, 54, 49, 52, 54, 53, 33, 64, 35, 36, 37, 94, 42, 40, 41, 95, 41, 43, 95, 43, 124, 125, 34,  
63, 62, 58, 58, 101, 117, 114, 103, 57, 101, 119, 117, 114, 103, 48, 119, 101, 117, 116, 57, 114,  
48, 52, 117, 48, 57, 53, 116, 117, 51, 48, 57, 50, 117, 116, 48, 57, 51, 117, 116, 48, 119, 57, 101]",  
"c1": "cf3f70a7-7565-44b0-ae3c-83bec549ea8e:104",  
"c2": "",  
"c3": "103",  
"c4": "2021-06-25 17:51:53",  
"c5": "1624614713.201",  
"c6": "!@#%$90weurtg103",  
"c7": "10357.0",  
"c8": "1.2510357E7",  
"c9": "9874510357",  
"id": "104"  
  }  
]
```

```
"c1":"cf3f70a7-7565-44b0-ae3c-83bec549ea8e:104",
"c2": "",
"c3":"103",
"c4":"2021-06-25 17:51:53",
"c5":"1624614713.201",
"c6":"!@#%90weurtg103",
"c7":"10357.0",
"c8":"1.2510357E7",
"c9":"9874510357",
"id":"103"
}
],
"pkNames":[
  "id"
]
}
```

JSON-C

JSON-C is similar to JSON. The difference lies in the delete operation. JSON data is stored in old, and JSON-C is stored in data. Data of the timestamp type is converted into a character string in the format of yyyy-mm-dd hh:mm:ss.

For details, see [Table 4-64](#).

Table 4-64 JSON-C parameter description

Parameter	Description
mysqlType	Field name and type in the source table.
id	Sequence number of an event operation defined in DRS. The value increases monotonically.
es	The time when the record is generated in the source database. The value is a 13-digit Unix timestamp in milliseconds.
ts	The time when the data is written to the target Kafka. The value is a 13-digit Unix timestamp in milliseconds.
database	Database name. For the Oracle database, set this parameter to schema .
table	Table name.
type	Operation type, such as DELETE, UPDATE, INSERT, and DDL.
isDdl	Whether the operation is a DDL operation.
sql	A DDL-defined SQL statement. The value is "".
sqlType	JDBC type of the fields in the source table.
data	Latest data, which is a JSON array. If type is set to INSERT , this parameter indicates the latest inserted data. If type is set to UPDATE , this parameter indicates the latest updated data. If type is set to DELETE , this parameter indicates the deleted data.

Parameter	Description
old	Old data. If type is set to UPDATE , the value indicates the data before update. If type is set to INSERT , the value is null .
pkNames	Primary key name

Common Escape Characters in JSON

Table 4-65 Escape Character

Character	Escape character
<	\u003d
>	\u003e
&	\u0026amp;
=	\u003d
'	\u0027

4.7.2 Kafka Authentication

PLAINTEXT

No security authentication mode is available. You only need to enter the IP address and port for connection.

Figure 4-2 PLAINTEXT

Destination Database

IP Address or Domain Name



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

Security Protocol

PLAINTEXT



Test Connection

SASL_PLAINTEXT

The SASL mechanism is used to connect to Kafka, and you need to configure SASL parameters.

Figure 4-3 SASL_PLAINTEXT

Destination Database

IP Address or Domain Name ?

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

Security Protocol ?

SASL Mechanisms ?

Username

Password

Table 4-66 Parameter settings

Parameter	Description
SASL Mechanisms	SASL is used by client. The following four items are supported. Kafka server uses the GSSAPI mechanism by default. <ul style="list-style-type: none"> ● GSSAPI ● PLAIN ● SCRAM-SHA-256 ● SCRAM-SHA-512
Token Delegation	Whether an agency token is used for authentication. This option is available when SCRAM-SHA-256 or SCRAM-SHA-512 is selected for SASL Mechanisms .
Username	Username for logging in to the database
Password	Password for the username

SSL

SSL is used to encrypt the connection to Kafka. Related parameters need to be configured.

Figure 4-4 SSL

Destination Database

IP Address or Domain Name ?

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

Security Protocol ?

Truststore Certificate

Truststore Certificate Password ?

Endpoint Identification Algorithm ?

Mutual SSL Authentication

Keystore Certificate

Truststore Certificate Password ?

Keystore Private Key Password ?

Table 4-67 Parameter settings

Parameter	Description
Truststore Certificate	SSL certificate with the file name extension .jks.
Truststore Certificate Password	Password of the certificate
Endpoint Identification Algorithm	Endpoint identification algorithm for verifying the host name of the server using the server certificate. This parameter is optional. If this parameter is left blank, host name verification is disabled.
Mutual SSL Authentication	Mutual SSL Authentication
Keystore Certificate	If mutual SSL authentication is enabled, you need to upload the mutual SSL authentication certificate with the file name extension .jks.
Keystore Certificate Password	Password of the mutual SSL authentication certificate. This option is available if mutual SSL authentication is enabled.
Keystore Private Key Password	(Optional) Password of the private key in the keystore certificate.

SASL_SSL

If the SASL and SSL are used, configure SSL and SASL parameters. For details, see [SASL_PLAINTEXT](#) and [SSL](#).

Figure 4-5 SASL_SSL

Destination Database

IP Address or Domain Name	<input type="text"/>	?
Ensure that the entered addresses belong to the same DB instance.		
Security Protocol	<input type="text" value="SASL_SSL"/>	?
SASL Mechanisms	<input type="text" value="GSSAPI"/>	?
Username	<input type="text"/>	
Password	<input type="text"/>	
Truststore Certificate	<input type="text"/>	Select
Truststore Certificate Password	<input type="text"/>	?
Endpoint Identification Algorithm	<input type="text"/>	?
Mutual SSL Authentication	<input checked="" type="checkbox"/>	
Keystore Certificate	<input type="text"/>	Select
Truststore Certificate Password	<input type="text"/>	?
Keystore Private Key Password	<input checked="" type="checkbox"/> <input type="text"/>	?
<input type="button" value="Test Connection"/>		

4.7.3 Forcibly Stopping Synchronization of PostgreSQL

This section describes how to clear the logical replication slot of the source database, how to synchronize sequence values, and how to reset the sequence values in the destination database when the source database cannot be connected after the PostgreSQL synchronization task is forcibly stopped.

Clearing the Logical Replication Slot of the Source Database

- Step 1** Log in to the source database as the source database user used in the synchronization task.
- Step 2** Query the name of the streaming replication slot of the database object selected in the synchronization task.

```
select slot_name from pg_replication_slots where database = 'database';
```

NOTICE

In the preceding command, *database* indicates the database selected in the synchronization task.

Step 3 Run the following statement to delete the streaming replication slot:

```
select * from pg_drop_replication_slot('slot_name');
```

NOTICE

In the preceding command, *slot_name* indicates the name of the streaming replication slot queried in [Step 2](#).

Step 4 Run the following statement to check whether the streaming replication slot is successfully deleted:

```
select slot_name from pg_replication_slots where slot_name = 'slot_name';
```

If the query result is empty, the streaming replication slot is deleted.

----End

Synchronizing Sequence Values

If sequence objects are not synchronized or the destination database is GaussDB, skip this section.

Step 1 Use a high-privilege account (with the USAGE permission for all sequences) to connect to the source database and run the following statement:

```
select 'SELECT pg_catalog.setval('||quote_literal(quote_ident(n.nspname))||','||quote_ident(c.relname))||', '||nextval(c.oid)||');' as sqls from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname <> 'information_schema' and not (c.relname='hwdrs_ddl_info_id_seq' and n.nspname='public') order by n.nspname, c.relname;
```

The query result is the SQL statement that needs to be executed in the destination database.

Step 2 Log in to the destination database as the destination database user used in the synchronization task and run the SQL statement queried in [Step 1](#) in the destination database.

Step 3 Run the following statement in the destination database to check the sequence value synchronization result:

```
SELECT n.nspname, c.relname, nextval(c.oid) from pg_class c join pg_namespace n on c.relnamespace=n.oid where c.relkind = 'S' and n.nspname !~'^pg_' and n.nspname <> 'information_schema' order by 1,2;
```

----End

Resetting Sequence Values in the Destination Database

If the source database failed and cannot be connected, you can reset the sequence values related to automatic increment or decrement columns in the destination database. If the source database can be connected, skip this section.

Step 1 Log in to the destination database as the destination database user used in the synchronization task.

- Step 2** Run the following statement to query the SQL statement for resetting the sequence value corresponding to the sequence that uses nextval as the default value of the table column:

```
set search_path to ''; select 'SELECT pg_catalog.setval('||quote_literal(quote_ident(s.sequence_schema))||'|' ||
quote_ident(s.sequence_name))||', (SELECT '||case when s.increment::int<0 then 'min(' else 'max(' end||
quote_ident(c.column_name))||')' ||case when s.increment::int<0 then '-1' else '+1' end||' FROM ' ||
quote_ident(c.table_schema))||'|' ||quote_ident(c.table_name))||');' as sqls from information_schema.columns c
join information_schema.sequences s on (position(quote_literal(quote_ident(s.sequence_schema))||'|' ||
quote_ident(s.sequence_name))||':regclass' in c.column_default) > 0) where c.data_type in ('bigint', 'int',
'integer', 'smallint', 'numeric', 'real', 'double precision', 'double') and c.column_default like 'nextval(%%%'
order by s.sequence_schema, s.sequence_name;
```

The query result is the SQL statement that needs to be executed in the destination database.

- Step 3** If the source database version is earlier than 10.0, skip this step. If the source database version is 10.0 or later, run the following statement in the destination database to query the SQL statement for resetting the sequence value corresponding to the additional column of the table identity column:

```
set search_path to ''; select 'SELECT pg_catalog.setval('||quote_literal(seqname))||', (SELECT '||case when
increment::int<0 then 'min(' else 'max(' end||colname||')' ||case when increment::int<0 then '-1' else '+1' end||'
FROM ' ||tablename||');' as sqls from (select objid::regclass::text, refobjid::regclass::text,
(pg_identify_object(refclassid,refobjid,refobjsubid)).identity, (pg_sequence_parameters(objid)).increment
from pg_depend where deptype='i' and refobjsubid>0 and objid in (select c.oid from pg_class c join
pg_namespace n on c.relnamespace=n.oid where c.relkind='S' and n.nspname !~ '^pg_' and
n.nspname<>'information_schema')) p(seqname,tablename,colname,increment);
```

The query result is the SQL statement that needs to be executed in the destination database.

- Step 4** Run the SQL statements queried in [Step 2](#) and [Step 3](#) in the destination database.

- Step 5** Run the following statement in the destination database to check the sequence value synchronization result:

```
SELECT n.nspname, c.relname, nextval(c.oid) from pg_class c join pg_namespace n on c.relnamespace=n.oid
where c.relkind = 'S' and n.nspname !~ '^pg_' and n.nspname<>'information_schema' order by 1,2;
```

----End

4.7.4 Creating Triggers and Functions to Implement Incremental DDL Synchronization for PostgreSQL

This section describes how to perform real-time synchronization from PostgreSQL to RDS PostgreSQL. You can create triggers and functions in the source database to obtain the DDL information of the source database, and then synchronize DDL operations to the destination database during the incremental synchronization phase.

Prerequisites

- The following DDL operations are supported:
 - Table-level synchronization: TRUNCATE (only for PostgreSQL 11 or later), DROP TABLE, and ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN, ADD CONSTRAINT, DROP CONSTRAINT and RENAME)
 - Database-level synchronization: TRUNCATE (only for PostgreSQL 11 or later), CREATE SCHEMA/TABLE, DROP TABLE, ALTER TABLE (including ADD COLUMN, DROP COLUMN, ALTER COLUMN, RENAME COLUMN,

ADD CONSTRAINT, DROP CONSTRAINT and RENAME), CREATE SEQUENCE, DROP SEQUENCE, ALTER SEQUENCE, CREATE INDEX, ALTER INDEX, DROP INDEX, CREATE VIEW, and ALTER VIEW

CAUTION

- Table-level synchronization: If data is inserted into a renamed table, the data will not be synchronized to the destination database.
 - Database-level synchronization: Tables that are created not using the CREATE TABLE statement in the source database will not be synchronized to the destination database. For example, you run CREATE TABLE AS to create a table or call a function to create a table.
 - DDL statements starting with comments cannot be synchronized and are ignored.
 - DDL statements executed in functions and stored procedures cannot be synchronized and are ignored.
-
- If the source and destination databases are of different versions, use SQL statements that are compatible with both the source and destination databases to perform DDL operations. For example, if the source database is PostgreSQL 11 and the destination database is PostgreSQL 12, run the following statement to change the column type from char to int:

```
alter table tablename alter column columnname type int USING columnname::int;
```
 - Check whether a table named **hwdrs_ddl_info**, a function named **hwdrs_ddl_function()**, and a trigger named **hwdrs_ddl_event** exist in the source database in public mode. If they exist, delete them.
 - During database-level synchronization, if a table without a primary key is created, run the following command to set the replication attribute of the table without a primary key to full.

```
alter table tablename replica identity full;
```

Procedure

NOTE

If the source is an RDS for PostgreSQL instance on the current cloud, you can create related objects as user **root**. If the "Must be superuser to create an event trigger" error is reported, you can submit a service ticket. For details about permissions of user **root** of RDS for PostgreSQL on the current cloud, see Relational Database Service User Guide.

Step 1 Connect to the database to be synchronized as a user who has permission to create event triggers.

Step 2 Run the following statements to create a table for storing DDL information:

```
DROP TABLE IF EXISTS public.hwdrs_ddl_info;
DROP SEQUENCE IF EXISTS public.hwdrs_ddl_info_id_seq;
CREATE TABLE public.hwdrs_ddl_info(
  id          bigserial primary key,
  ddl        text,
  username   varchar(64) default current_user,
  txid       varchar(16) default txid_current()::varchar(16),
  tag        varchar(64),
  database   varchar(64) default current_database(),
  schema     varchar(64) default current_schema,
```

```
client_address    varchar(64) default inet_client_addr(),
client_port       integer default inet_client_port(),
event_time       timestamp default current_timestamp
);
```

Step 3 Run the following statements to create a function:

```
CREATE OR REPLACE FUNCTION public.hwdrs_ddl_function()
  RETURNS event_trigger
  LANGUAGE plpgsql
  SECURITY INVOKER
AS $BODY$
declare ddl text;
declare real_num int;
declare max_num int := 50000;
begin
  if (tg_tag in ('CREATE TABLE','ALTER TABLE','DROP TABLE','CREATE SCHEMA','CREATE SEQUENCE','ALTER SEQUENCE','DROP SEQUENCE','CREATE VIEW','ALTER VIEW','DROP VIEW','CREATE INDEX','ALTER INDEX','DROP INDEX')) then
    select current_query() into ddl;
    insert into public.hwdrs_ddl_info(ddl, username, txid, tag, database, schema, client_address, client_port, event_time)
      values (ddl, current_user, cast(txid_current() as varchar(16)), tg_tag, current_database(), current_schema, inet_client_addr(), inet_client_port(), current_timestamp);
    select count(id) into real_num from public.hwdrs_ddl_info;
    if real_num > max_num then
      if current_setting('server_version_num')::int<100000 then
        delete from public.hwdrs_ddl_info where id<(select min(id)+1000 from public.hwdrs_ddl_info) and
          not exists (select 0 from pg_locks l join pg_database d on l.database=d.oid where
            d.datname=current_catalog and pid<>pg_backend_pid() and locktype='relation' and
            relation=to_regclass('public.hwdrs_ddl_info_pkey')::oid and mode='RowExclusiveLock');
      else
        delete from public.hwdrs_ddl_info where id<(select min(id)+1000 from public.hwdrs_ddl_info) and
          (xmax=0 or coalesce(txid_status(xmax::text::bigint), '')<>'in progress');
      end if;
    end if;
  end if;
end;
$BODY$;
```

Step 4 Run the following statements to grant necessary permissions to the objects created in [Step 2](#) and [Step 3](#):

```
GRANT USAGE ON SCHEMA public TO public;
GRANT SELECT,INSERT,DELETE ON TABLE public.hwdrs_ddl_info TO public;
GRANT SELECT,USAGE ON SEQUENCE public.hwdrs_ddl_info_id_seq TO public;
GRANT EXECUTE ON FUNCTION public.hwdrs_ddl_function() TO public;
```

Step 5 Run the following statement to create a DDL event trigger:

```
CREATE EVENT TRIGGER hwdrs_ddl_event ON ddl_command_end EXECUTE PROCEDURE
public.hwdrs_ddl_function();
```

Step 6 Run the following statement to set the created event trigger to enable:

```
ALTER EVENT TRIGGER hwdrs_ddl_event ENABLE ALWAYS;
```

Step 7 Return to the DRS console and create a PostgreSQL to RDS PostgreSQL synchronization task.**Step 8** After the synchronization task is complete, run the following statements to delete the created tables, functions, and triggers.

```
DROP EVENT trigger hwdrs_ddl_event;
DROP FUNCTION public.hwdrs_ddl_function();
DROP TABLE public.hwdrs_ddl_info;
```

----End

5 FAQs

5.1 Product Consulting

5.1.1 What Are Regions and AZs?


Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center. Each region is completely independent, improving fault tolerance and stability. After a resource is created, its region cannot be changed.
- An AZ is a physical location using independent power supplies and networks. Faults in an AZ do not affect other AZs. A region can contain multiple AZs, which are physically isolated but interconnected through internal networks. This ensures the independence of AZs and provides low-cost and low-latency network connections.

Figure 5-1 shows the relationship between regions and AZs.

Figure 5-1 Region and AZ



Selecting a Region

You are advised to select a region close to you or your target users. This reduces network latency and improves access rate.

Selecting an AZ

When determining whether to deploy resources in the same AZ, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs in the same region.
- For low network latency, deploy resources in the same AZ.

Regions and Endpoints

Before using an API to call resources, specify its region and endpoint. For more details, see [Regions and Endpoints](#).

5.1.2 What Is DRS?

DRS is a stable, efficient, and easy-to-use cloud service for database migration and synchronization in real time.

It simplifies data migration processes and reduces migration costs.

You can use DRS to quickly transmit data between different DB engines.

Real-Time Migration

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

You can perform a migration task over multiple types of networks, such as public networks, VPCs, VPNs, and direct connections. With these network connections, you can migrate between different cloud platforms, from on-premises databases to cloud databases, or between cloud databases across regions.

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

Backup Migration

It often becomes necessary to hide the 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.

You can use backup migration when you want to migrate on-premises databases to the cloud.

Without connecting to your sources, DRS can help you complete data migration.

Real-Time Synchronization

Real-time synchronization refers to the real-time flow of key service data from sources to destinations through a synchronization instance while consistency of data can be ensured.

It is different from migration. Migration means moving your overall database from one platform to another. Synchronization refers to the continuous flow of data between different services.

You can use real-time synchronization in many scenarios such as real-time analysis, report system, and data warehouse environment.

Real-time synchronization is mainly used for synchronizing tables and data. It can meet various requirements, such as many-to-one, one-to-many synchronization, dynamic addition and deletion of tables, and synchronization between tables with different names.

5.1.3 Can DRS Migrate RDS Primary/Standby Instances?

Yes. DRS provides high availability and can migrate a single RDS instance or RDS primary/standby instances. DRS can automatically rebuild the databases connection after a short interruption and resumes data transfer from the point when the connection was lost to ensure the continuity and consistency of data synchronization.

If the HA design of the source database meets the requirements of floating IP address connections and RPO is 0 during a switchover, DRS supports migration of primary/standby instances without manual intervention.

If the HA design does not meet the requirements of floating IP address connections and RPO is 0 during a switchover, the following situations may occur:

- The floating IP address is used and RPO may be 0 during a switchover. In this situation, the database can be connected, but DRS will identify data interruption (if data loss occurs during the switchover) and display a message indicating that the task fails. You can only reset the migration task.
- A fixed IP address is used and RPO is 0 during the switchover. In this situation, the migration is supported only when the instance is running properly.
- The floating IP address is used and zero RPO cannot be ensured during a switchover. In this situation, the database can be connected, but DRS will identify data interruption (if data loss occurs during the switchover) and display a message indicating that the task fails. You can only reset the migration task.

If the destination is primary/standby instances, DRS can ensure that the source data is completely migrated to the destination database. However, the switchover of the destination database cannot ensure zero RPO. As a result, data in the destination database may be incomplete.

5.1.4 Does DRS Support Resumable Uploads?

In database migration scenarios, if a migration task fails due to unavoidable problems (such as network fluctuation), DRS records the current parsing and

replay point (which is the basis of database internal consistency) and then resumes data transfer from the point to ensure data integrity.

For incremental migration, DRS automatically retries for multiple times. For full migration of MySQL databases, the system automatically resumes the migration for three times by default. After the number of automatic retry failures reaches a specified value, the task becomes abnormal. You need to analyze the cause based on logs and try to rectify the blocking point (for example, the database password is changed). If the environment cannot be restored and the required logs have been eliminated, you can use the reset the task.

5.1.5 What Are the Differences Between Real-Time Migration and Real-Time Synchronization?

Item	Real-Time Migration	Real-Time Synchronization
Scenario	Migration can be performed between different cloud platforms, from on-premises databases to cloud databases, or on cloud databases across regions.	Real-time analysis, report system, and data warehouse environment
Characteristics	Homogeneous databases are migrated as a whole. Tables, data, indexes, views, stored procedures, functions, database accounts, and database parameters can be migrated at the table level, database level, or all dimensions.	Maintains continuous data flow between different services, synchronizes tables and data, and meets various flexibility requirements. Objects can be migrated at the table level or database level. Data synchronization between heterogeneous databases is supported.
Functions and features	<ul style="list-style-type: none"> • Users can be migrated. • Parameters can be compared. 	<ul style="list-style-type: none"> • Synchronization objects can be edited in the incremental phase, and added or deleted dynamically. • Object names mapped to the destination database can be changed, so the names of tables and databases in the source and destination databases can be different. • Data processing is supported. You can add rules, such as data filtering and column processing, for selected objects.
Remarks	Different data flows support different functions and features. For details, see Precautions .	Different data flows support different functions and features. For details, see Precautions .

5.1.6 How Do I Solve the Table Bloat Issue?

In the full migration phase, DRS uses the row-level parallel migration mode to ensure migration performance and transmission stability. If the source database data is compact, table bloat may occur after data is migrated to the RDS MySQL database. As a result, the disk space required is much greater than that of the source database. In this case, you can run the following command in the destination database to free up the space:

```
optimize table table_name
```

NOTE

The OPTIMIZE TABLE command locks tables. Do not run this command when you operate table data. Otherwise, services may be affected.

5.1.7 How Do I Select RDS Read Replicas on the DRS Console?

RDS read replicas cannot be selected on the DRS console. You can select **Self-built on ECS** and enter the read replica IP address and port number to connect to the instance.

5.1.8 How Does DRS Affect the Source and Destination Databases?

Impact on the Source

- During the initialization of a full migration or synchronization task, DRS needs to query all inventory data in the source database. DRS uses simple SQL statements to query data, and the query speed is limited by the I/O performance and network bandwidth of the source database. Generally, if the bandwidth is not limited, the query workload of the source database will be increased by 50 MB/s and 2 to 4 vCPUs will be occupied. If the source database is read concurrently, about 6 to 10 sessions are occupied.
 - Fewer than eight sessions are used to query some system tables, such as tables, views, and columns in the information_schema database, in the source database.
 - Fewer than four sessions are used to query shards in the source database. For example, in the following statement, the conditions following **select** and **where** contain only the primary key or unique key.

```
select id from xxx where id>12345544 and limit 10000,1;
```
 - Fewer than four sessions are used to query SQL statements. For example, in the following statement, the information after **select** is all column names in the table, and the condition after **where** contains only the primary key or unique key.

```
select id,name,msg from xxx where id>12345544 and id<=12445544;
```
 - The SQL statement for locking a table without a primary key is similar to the following statement. The table is locked to obtain the consistency point of the table without a primary key. After the table is locked, a connection is obtained to unlock the table.

```
flush table xxx with read lock
```

```
lock table xxx read
```

- In the incremental phase, there is no pressure on the source database. Only one dump connection is available to listen to binlog incremental data in real time.

Impact on the Destination Database

- During the initialization of a full migration or synchronization task, DRS needs to write structures, inventory data, and indexes of the source database to the destination database in sequence. Generally, the total number of sessions is less than 16.
 - Fewer than eight sessions are used to create structures.
 - Fewer than eight sessions are writing data. Example:

```
insert into xxx (id,name,msg) values (xxx);
```
 - Fewer than eight sessions are used to create indexes. Example:

```
alter table xxx add index xxx;
```
- In the incremental phase, DRS parses the incremental data in the binlog file of the source database into SQL statements and executes the SQL statements in the destination database. Generally, the total number of sessions is less than 64.
 - DDL statements are executed in serial mode. When a DDL statement is executed, no other DML statement is executed.
 - There are a maximum of 64 DML connections (short connections with a timeout interval of 30 seconds). The DML statements include insert, update, delete, and replace.

NOTE

To evaluate the impact on the source database, you can create a test task and adjust the migration policy by using rate limiting or run the test during off-peak hours.

5.1.9 Can DRS Migrates Table Structures Only?

DRS is a cloud service used for real-time data transfer. Currently, DRS cannot migrate table structures only but not data. For details about the objects supported by each data flow, see the following links.

5.2 Network and Security

5.2.1 What Security Protection Policies Does DRS Have?

Network

- Uses security groups to ensure that the sources of access are trusted.
- Uses SSL channels to encrypt data during transmission.

5.2.2 What Can I Do If the Network Is Disconnected During the Migration?

If the network is disconnected during the migration, you can view the task status first. If a full or incremental task fails, click **Resume** in the **Operation** column.

Full migration

Incremental migration

5.2.3 How Can I Configure a VPC Security Group to Allow Traffic from an EIP?

By default, a VPC on the current cloud is isolated from external networks for security reasons. You cannot use an EIP outside a VPC (for example, an EIP of another cloud database or an on-premise database) to access DB instances inside the VPC. However, the replication instance or destination database in a VPC needs to connect to an external EIP to migrate data.

Therefore, you need to add an outbound rule to a security group to allow access from specific external EIPs and ports outside the VPC. The outbound rule allows the replication instance EIP to access the destination database EIP.

5.2.4 What Can I Do If the Network Connection Between the Replication Instance and Database Is Abnormal?

Before data migration, ensure that network preparations and security rule settings are complete. If the connection is abnormal, check whether the network configuration is correct.

This section uses the migration from MySQL to RDS MySQL as an example to describe three migration scenarios: cross-cloud real-time migration, on-premises database migration, and real-time migration of self-built ECS databases.

Cross-Cloud Real-Time Migration

1. Network settings

Enable public accessibility for the source database.

- Source database network settings:

Enable public accessibility for the source database.

For details, see related documents provided by Alibaba Cloud.

- 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.

2. Security rules

- Source database security group settings:

Add the EIP of the replication instance to the whitelist of the source MySQL DB instance to allow the access from the EIP.

Before configuring the whitelist for the source database, obtain the EIP of the DRS replication instance. You can find the EIP on the **Configure Source and Destination Databases** page after creating the replication instance on the DRS console.

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 you must ensure that the above does not pose a risk to your services.

After the migration is complete, you can delete the configuration from the whitelist.

- 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. DRS can directly write data to the destination database.
 - Configure the security group of the VPC where the destination database is located to ensure that the IP addresses and listening ports of the DRS instance are allowed to access the on-premises database.

Real-Time Migration of On-Premises Databases

1. Network settings

- Source database network settings:

You can migrate on-premises MySQL databases to the RDS MySQL databases 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 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 attempts to access the destination database through a VPN, ensure that the VPN service is enabled and the source database can communicate with the destination RDS MySQL database.
- If the source database attempts to access the destination database through a public network, you do not need to configure the destination RDS MySQL database.

2. Security rules

- a. Source database security group settings:

- If the migration is performed over a public network, add the EIP 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. Before setting the network whitelist, obtain the EIP of the replication instance.

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

- If the migration is performed over a VPN network, add the private IP address of the DRS migration 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.

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

- 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. DRS can directly write data to the destination database.
 - Configure the security group of the VPC where the destination database is located to ensure that the IP addresses and listening ports of the DRS instance are allowed to access the on-premises database.

Real-Time Migration of Self-Built Databases on the ECS

1. 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 the same VPC, the networks are interconnected by default.
 - 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.
2. 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.

Checking iptables Settings

If the source database is a self-built database on an ECS and cannot be connected after the preceding operations are performed, check the iptables settings. If the DRS frequently initiates connection requests and fails, the HOSTGUARD service adds the requested IP address to the blacklist.

1. Log in to the ECS.
2. Run the following command to check whether any DENY-related project contains the IP address of the DRS instance. The project name is **IN_HIDS_MYSQLD_DENY_DROP**.
iptables --list
3. If yes, run the following command to query the iptables inbound rule list and obtain the rule ID (line-numbers):

iptables -L INPUT --line-numbers

4. Run the following command to delete the inbound rules that deny the IP address of the DRS instance: (Note: Delete the rules from the end to the beginning. Otherwise, line-numbers will be updated and you need to query again.)

iptables -D *Project_name Rule_ID*

5. Delete the iptables rules and test the connection again.

5.3 Permissions Management

5.3.1 Which MySQL Permissions Are Required for DRS?

DRS has certain permission requirements on accounts during migration. This section describes the permission requirements on the MySQL engine.

Permission

- You must have the login permission of the source and destination database connection accounts. If you do not have the account, perform the following operations to create one. user1 is used as an example.
Reference statement: **CREATE USER 'user1'@'host' IDENTIFIED BY 'password'**
- [Table 5-1](#) lists the required permissions for real-time migration.

Table 5-1 Permission requirements and reference statements

Function Modules	Source/Service Database	Destination/DR Database
Real-time migration	<p>Full migration: SELECT, SHOW VIEW, and EVENT</p> <p>Reference statement: GRANT SELECT, SHOW VIEW, EVENT ON *.* TO 'user1';</p> <p>Full+incremental migration: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT</p> <ul style="list-style-type: none"> • REPLICATION SLAVE and REPLICATION CLIENT are global permissions and must be enabled separately. The reference statement is as follows: GRANT REPLICATION SLAVE, REPLICATION CLIENT ON *.* TO 'user1'; • SELECT, SHOW VIEW, EVENT, and LOCK TABLES are non-global permissions. The reference statement is as follows: GRANT SELECT, SHOW VIEW, EVENT, LOCK TABLES, ON [Database to be migrated].* TO 'user1'; 	<p>Full migration: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</p> <p>Reference statement: GRANT SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER ON *.* TO 'user1' WITH GRANT OPTION;</p> <p>Full+incremental migration: SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES, and WITH GRANT OPTION. If the destination database version is in the range 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</p> <p>Reference statement: GRANT SELECT, CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, REFERENCES ON [Databases to be migrated].* TO 'user1' WITH GRANT OPTION;</p>

 **NOTE**

Run **flush privileges;** after executing the preceding reference statements. Make the authorization take effect.

- Account migration:
The user must have the SELECT permission of mysql.user.
Reference statement:

```
GRANT SELECT ON mysql.user TO 'user1'@'host' ;
```

```
GRANT SELECT ON mysql.user_view TO 'user1';
```

The destination database users must have the SELECT, INSERT, UPDATE, DELETE, and WITH GRANT OPTION permissions on all databases.

Reference statement: **GRANT SELECT, INSERT, UPDATE, DELETE ON *.* TO 'user1' WITH GRANT OPTION**

Actions

- Create a user.

Operation:

```
CREATE USER 'username'@'host' IDENTIFIED BY 'password';
```

· **username**: indicates the account to be created.

· **host**: indicates the host that allows the account to log in. If the account is allowed to log in to the database from any host, use %.

· **password**: indicates the password of the account.

- Grant corresponding permissions.

Operation:

```
GRANT privileges ON databasename.tablename TO 'username'@'host' WITH GRANT OPTION;
```

```
flush privileges;
```

· **privileges**: indicates the operation permissions granted to the account, such as SELECT, INSERT, and UPDATE. To grant all permissions to the account, use ALL.

· **databasename**: indicates the database name. To grant the account with all database operation permissions, use *.

· **tablename**: indicates table name. To grant the account with all table operation permissions, use *.

· **username**: indicates the account to be authorized.

· **host**: indicates the host that allows the account to log in. If the account is allowed to log in from any host, use %.

· **WITH GRANT OPTION**: indicates that the permission to use the GRANT command is granted to the account. This parameter is optional.

5.3.2 How Can I Import Users and Permissions from the Source to the Destination Database?

Step 1 Log in to an ECS that can access the source database.

Step 2 Run the following command, enter the password as prompted, and press **Enter** to export the source database users to the **users.sql** temporary file:

```
mysql -h 'host' -u 'user' -p -N $@ -e "SELECT CONCAT('SHOW GRANTS FOR ''', user, ''@'', host, '';) AS query FROM mysql.user" > /tmp/users.sql
```

host indicates the IP address of the source database and **user** indicates the username of the source database.

Step 3 Run the following command to export the authorization information of the users from the source database to the **grants.sql** file:

```
mysql -h 'host' -u 'user' -p -N $@ -e "source /tmp/users.sql" > /tmp/grants.sql  
sed -i 's/$/;/g' /tmp/grants.sql
```

host indicates the IP address of the source database and *user* indicates the username of the source database.

Step 4 After the preceding command has been executed successfully, open the **grants.sql** file. Information similar to the following is displayed:

```
-- Grants for root@%  
GRANT ALL PRIVILEGES ON *.* TO 'root'@'%';  
  
-- Grants for testt@%  
GRANT SELECT, INSERT, UPDATE, DELETE ON *.* TO 'testt'@'%';  
  
-- Grants for debian-sys-maint@localhost  
GRANT ALL PRIVILEGES ON *.* TO 'debian-sys-maint'@'localhost' WITH GRANT OPTION;  
  
-- Grants for mysql.session@localhost  
GRANT SUPER ON *.* TO 'mysql.session'@'localhost';  
GRANT SELECT ON `performance_schema`.* TO 'mysql.session'@'localhost';  
GRANT SELECT ON `mysql`.`user` TO 'mysql.session'@'localhost';  
  
-- Grants for mysql.sys@localhost  
GRANT USAGE ON *.* TO 'mysql.sys'@'localhost';  
GRANT TRIGGER ON `sys`.* TO 'mysql.sys'@'localhost';  
GRANT SELECT ON `sys`.`sys_config` TO 'mysql.sys'@'localhost';  
  
-- Grants for root@localhost  
GRANT ALL PRIVILEGES ON *.* TO 'root'@'localhost' WITH GRANT OPTION;  
GRANT PROXY ON "" TO 'root'@'localhost' WITH GRANT OPTION;
```

Step 5 The information displayed in [Step 4](#) shows all users of the source database and their permissions. Add the required users one by one to the RDS MySQL database on the current cloud.

----End

5.4 Real-Time Migration

5.4.1 When Can I Stop a Migration Task?

You can refer to the following methods to check whether the task can be stopped. Before stopping the task, ensure that:

1. At least one complete data comparison is performed during off-peak hours.
2. Service cutover is completed.
 - a. Interrupt services first. If the workload is not heavy, you may do not need to interrupt the services.
 - b. Run the following statement on the source database (MySQL is used as an example) and check whether there are statements executed by new sessions within 1 to 5 minutes. If not, the service is 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.

- c. When the real-time synchronization delay is 0s and remains stable for a period, you can perform a data-level comparison between the source and destination databases. For details about the time required, refer to the comparison 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.
 - d. Determine a proper time to cut the services over to the destination database. Then, services can be used externally again.
3. Stopping a task only deletes the replication instance, and the migration task is still in the task list. You can choose whether to delete the task.

5.4.2 How Do I Maintain the Original Service User Permission System After Definer Is Forcibly Converted During MySQL Migration?

Definer is used in views, stored procedures, triggers, and events. Definer does not restrict the permission to invoke objects, instead the permission to access the database. If you select **Yes** for **Migrate Definer to User** during MySQL migration, the Definers of all source database objects will be migrated to the user. The user continues to use the original services without authorization. (Users, permissions, and passwords are migrated). Other users do not have permissions on database objects unless these users are authorized.

The following procedures describe how to use database commands to authorize users.

- Step 1** Ensure that the new user (Definer uses the specified account) has sufficient permission to execute view- and stored procedure-related SQL statements.
- Step 2** Log in to the destination database using the MySQL official client or other tools.
- Step 3** Run the following command to view details about permissions of the user to be authorized:

```
show grants for 'user'@'host';
```
- Step 4** To ensure that the original service does not report an error, run the following command to grant the user the operation permissions the involved database objects do not have:

```
grant select,insert,update,delete on db_name.* to 'user'@'host';
```

Generally, the permissions to access the database are as follows: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, and EXECUTE. You need to check the permissions that are missing based on the database object, and then perform the authorization operation.

For stored procedures and functions, ensure that the user has the EXECUTE permission. The authorization command is as follows:

```
grant execute on db_name.function_name to 'user'@'host';
```

- Step 5** Use the authorized account to access the destination database. If the access is successful, the authorization is successful. Note: If the following information is displayed when a stored procedure or function is invoked in a Java project, the **mysql.proc** database must be authorized: `Java.sql.SQLException: User does not have access to metadata required to determine stored procedure parameter types.` If rights can not be granted, configure connection with `"noAccessToProcedureBodies=true"` to have driver generate parameters that represent INOUT strings irregardless of actual parametertypes
- ```
grant select on mysql.proc to 'user'@'host';
```

----End

### 5.4.3 What Can I Do If the Invoking Permission Problem Occurs After the MySQL Stored Procedure Is Migrated to the Cloud?

After the MySQL stored procedure is migrated to the cloud, an error may occur when the stored procedure or function is invoked due to permission problems.

The method varies with Definer policies. This section uses user1 as an example to describe how to solve this problem in two Definer policies.

#### Policy 1

On the **Destination Database** page, enter the database username **user1**, and select **OK** for **Migrate Definer to User**.

In this policy, after the Definers of all stored procedures and methods in the source database are migrated to the destination database, the account is automatically changed to user1, and the value of host is automatically changed to %. If a stored procedure fails to be invoked in the destination database, perform the following operations:

- Step 1** Log in to the RDS MySQL DB instance of the destination database as the user1.
- Step 2** Grant the execute permission to the account that you want to use to invoke a stored procedure.
- Step 3** Run the following statement to use user1 to grant other accounts the permission to execute stored procedures:

**user** indicates other accounts that need to invoke the stored procedure.

```
GRANT EXECUTE ON db.* TO user;
```

- Step 4** To invoke a stored procedure using Java, run the following statement to use user1 to grant other accounts the permission to query the **mysql.proc** table:

The following is the authorization statement, in which **user** indicates the account that needs to invoke the stored procedure:

```
GRANT SELECT ON mysql.proc TO 'user'@'%';
```

----End

## Policy 2

On the **Destination Database** page, enter the database username **user1**, and select **Cancel** for **Migrate Definer to User**.

In this policy, the account and host in the source database remain unchanged after the Definers of all stored procedures and methods are migrated to the destination database. You need to migrate all users in the source database by referring to . In this way, the permission system of the source database remains unchanged.

If you do not migrate account permissions or some accounts cannot be migrated, you are advised to use [Policy 1](#).

### 5.4.4 How Do I Ensure that All Services on the Database Are Stopped?

To ensure that all services on the database are stopped, perform the following steps:

- Step 1** Run the following statement on the source database to check whether active connections exist:

```
show processlist;
```

Figure 5-2 Checking active connections

```
mysql> show processlist;
+----+-----+-----+-----+-----+-----+-----+-----+
| Id | User | Host | db | Command | Time | State | Info |
+----+-----+-----+-----+-----+-----+-----+-----+
| 3802808 | root | 192.168.0.117:41686 | NULL | Binlog Dump GTID | 58680 | Master has sent all binlog to slave; waiting for more updates | NULL |
| 3866358 | root | 192.168.0.64:55546 | NULL | Query | 0 | starting | show processlist |
+----+-----+-----+-----+-----+-----+-----+-----+
rows in set (0.00 sec)
```

- Step 2 Optional:** If there are active connections, locate the service processes based on the values in the **Host** column in the command output and stop the service processes.
- Step 3** Run the following statement in the source database to check the binlog position. Then, record the two values in the **file** and **position** columns as **ckpt1**:

```
show master status;
```

Figure 5-3 Viewing the binlog position

```
mysql> show master status;
+-----+-----+-----+-----+-----+
| File | Position | Binlog_Do_DB | Binlog_Ignore_DB | Executed_Gtid_Set |
+-----+-----+-----+-----+-----+
| mysql-bin.005290 | 197 | | | 67811045-de76-11e9-84c5-fa163e7a0434:1-592564543 |
+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

- Step 4** Wait for more than 30s. Run the following statement in the source database to check the binlog position again. Then, record the two values in the **file** and **position** columns as **ckpt2**. If **ckpt1** and **ckpt2** are equal, no data is written to the source database.

```
show master status;
```

----End



## 5.4.5 What Can I Do When Scheduled Task Failed to Start Leaving Message "can not get agency token"

When you use a subaccount to use scheduled task startup function, the **account entrustment** function must be used. Otherwise, the scheduled task fails to be started, leaving message "can not get agency token".

### Solution

Two solutions are provided as follows:

- Method 1: Use your primary account to create a task and select **Start at a specified time** for **Start Time**.
- Method 2: Use the primary account to add the Security Administrator permission to the user group to which the subaccount belongs. Then, create a task again, and select **Start at a specified time** for **Start Time**.
- Method 3: Create a task again and select **Start upon task creation** for **Start Time**.

## 5.4.6 What Can I Do If MyISAM Tables Are Not Supported by RDS for MySQL?

Currently, RDS for MySQL does not support the MyISAM engine due to the following reasons.

- MyISAM engine tables do not support transactions and support only table-level locks. As a result, read and write operations conflict with each other.
- MyISAM has a defect in protecting data integrity, which may cause database data damage or even data loss.
- If data is damaged, MyISAM does not support data restoration provided by RDS for MySQL and requires manual restoration.
- Data can be transparently migrated from MyISAM to InnoDB, which does not require code modification for tables.

During migration, DRS automatically converts MyISAM to InnoDB. The MyISAM engine table does not support transactions. To ensure data consistency of the MyISAM table, DRS uses primary keys to ensure final data consistency. If you need to migrate MyISAM tables without primary keys, you are advised to start the migration task when no service is running to ensure data consistency.

## 5.4.7 What Are the Precautions for Migrating Data from an Earlier Version MySQL to MySQL 8.0?

Based on MySQL 5.7, some new features have been added to MySQL 8.0. There are performance differences between the two versions. Before migration, you need to analyze compatibility and provide a corresponding solution. The following shows the analysis:

- Compatibility analysis  
MySQL 8.0 and MySQL 5.7 Community Edition are analyzed as follows:

- a. Compatibility does not affect migration, but the solutions are different.

| Compatibility           | Check Item                                                                                                                        | Function                 | Status    | Solution                                     |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------|--------------------------|-----------|----------------------------------------------|
| Data types or functions | ENCODE()                                                                                                                          | Encryption               | Deleted   | Replaced by AES_ENCRYPT()                    |
|                         | DECODE()                                                                                                                          | Decryption               | Deleted   | Replaced by AES_DECRYPT()                    |
|                         | ENCRYPT()                                                                                                                         | Encryption               | Deleted   | Replaced by SHA2()                           |
|                         | DES_ENCRYPT()                                                                                                                     | Encryption               | Deleted   | Replaced by AES_ENCRYPT()                    |
|                         | DES_DECRYPT()                                                                                                                     | Decryption               | Deleted   | Replaced by AES_DECRYPT()                    |
|                         | JSON_APPEND()                                                                                                                     | Adds JSON elements.      | Deleted   | Replaced by JSON_ARRAY_APPEND()              |
|                         | PASSWORD()                                                                                                                        | Changes a user password. | Deleted   | ALTER USER user IDENTIFIED BY 'auth_string'; |
|                         | JSON_MERGE()                                                                                                                      | Merges multiple JSONs.   | Discarded | Replaced by JSON_MERGE_PRESERVE()            |
| SQL MODE                | NO_AUTO_CREATE_USER, DB2, MAXDB, MSSQL, MYSQL323, MYSQL40, ORACLE, POSTGRESQL, NO_FIELD_OPTIONS, NO_KEY_OPTIONS, NO_TABLE_OPTIONS | -                        | Deleted   | -                                            |

| Compatibility                 | Check Item                                                | Function           | Status  | Solution                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------|-----------------------------------------------------------|--------------------|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Foreign key constraint length | The constraint name cannot be greater than 64 characters. | -                  | -       | <pre>SELECT TABLE_SCHEMA, TABLE_NAME FROM INFORMATION_SCHEMA.T ABLES WHERE TABLE_NAME IN (SELECT LEFT(SUBSTR(ID,INSTR(ID, '/')+1), INSTR(SUBSTR(ID,INSTR(I D,'/')+1),'_ibfk_')-1) FROM INFORMATION_SCHEMA.I NNODB_SYS_FOREIGN WHERE LENGTH(SUBSTR(ID,INST R(ID,'/')+1))&gt;64);</pre> <p>Use the ALTER TABLE statement to adjust the length.</p> |
| Features                      | Use the GRANT statement to create users.                  | -                  | Deleted | CREATE USER                                                                                                                                                                                                                                                                                                                                      |
|                               | Use the GRANT statement to modify user information.       | -                  | Deleted | ALTER USER                                                                                                                                                                                                                                                                                                                                       |
|                               | IDENTIFIED BY PASSWORD 'auth_string'                      | Sets new passwords | Deleted | IDENTIFIED WITH auth_plugin AS 'auth_string'                                                                                                                                                                                                                                                                                                     |
|                               | \N in a SQL statement                                     | NULL               | Deleted | Replaced by NULL                                                                                                                                                                                                                                                                                                                                 |

| Compatibility | Check Item                 | Function                                                                                  | Status  | Solution                                     |
|---------------|----------------------------|-------------------------------------------------------------------------------------------|---------|----------------------------------------------|
|               | PROCEDURE ANALYSE() syntax | Specifies the recommended field type is provided after the MySQL field value is analyzed. | Deleted | -                                            |
|               | Spatial functions          | -                                                                                         | -       | -                                            |
|               | mysql_install_db           | Initialization                                                                            | Deleted | mysqld --initialize or --initialize-insecure |

- b. The following items affect the migration. You need to check in advance.

| Compatibility      | Check Item                                                                                                                                                                                             | Function | Status | Solution                     | Original Usage                                                                                 |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|--------|------------------------------|------------------------------------------------------------------------------------------------|
| Reserving keywords | cume_dist, dense_rank, empty, except, first_value, grouping, groups, json_table, lag, last_value, lateral, lead, nth_value, ntile, of, over, percent_rank, rank, recursive, row_number, system, window | -        | Added  | SET sql_mode = 'ANSI_QUOTES' | Name: database, table, index, column, alias, view, stored procedure, partition, and tablespace |

| Com patib ility       | Check Item                                                                                | Fun ctio n                 | St at us     | Solution                                                                                                                                                                                                                                                                                                                | Original Usage                                     |
|-----------------------|-------------------------------------------------------------------------------------------|----------------------------|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------|
| Char acter set        | UTF8MB3                                                                                   | -                          | Di sc ar ded | Replaced by UTF8MB4.                                                                                                                                                                                                                                                                                                    | -                                                  |
| Partit ion table name | Partition tables of storage engines that do not support local partitions are not allowed. | -                          | D el et ed   | <p>SELECT TABLE_SCHEMA, TABLE_NAME FROM INFORMATION_SCHEMA.TABLES WHERE ENGINE NOT IN ('innodb', 'ndbcluster') AND CREATE_OPTIONS LIKE '%partitioned%';</p> <p>You can use either of the following methods:</p> <p>(1) ALTER TABLE table_name ENGINE=INNODB;</p> <p>(2) ALTER TABLE table_name REMOVE PARTITIONING;</p> | MyISAM is not supported.                           |
| Synta x               | group by... asc/desc                                                                      | Asc end ing/ Des cen din g | D el et ed   | Replaced by the ORDER By clause.                                                                                                                                                                                                                                                                                        | View and function                                  |
| Nam e lengt h         | The view name cannot be greater than 64 characters.                                       | -                          | -            | ALTER                                                                                                                                                                                                                                                                                                                   | The value can contain a maximum of 255 characters. |
|                       | The enum or set element contains a maximum of 255 characters.                             | -                          | -            | Handled by users.                                                                                                                                                                                                                                                                                                       | The value can contain a maximum of 64 KB.          |

| Compatibility                | Check Item                                                              | Function                                                      | Status | Solution                                                                                                                                                                                                                                                                                                                                                          | Original Usage |
|------------------------------|-------------------------------------------------------------------------|---------------------------------------------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Upper and lower case letters | lower_case_table_names                                                  | Specifies whether to set the MySQL table name case sensitive. | -      | <p>If this parameter is set to <b>1</b> during the upgrade, ensure that the schema and table names are in lowercase.</p> <pre>SELECT TABLE_NAME FROM INFORMATION_SCHEMA.TABLES WHERE TABLE_NAME != LOWER(TABLE_NAME) AND TABLE_TYPE = 'BASE TABLE';</pre> <pre>SELECT SCHEMA_NAME FROM INFORMATION_SCHEMA.SCHEMATA WHERE SCHEMA_NAME != LOWER(SCHEMA_NAME);</pre> | -              |
| Triggers                     | Check whether there is an empty definition or invalid creation context. | -                                                             | -      | Use the SHOW TRIGGERS statement to check the character_set_client, collation_connection, and Database Collation attributes.                                                                                                                                                                                                                                       | -              |

- Change the default value of the system variable.

The analysis of default values of MySQL 5.7 and MySQL 8.0 Community Edition shows that default values do not affect the migration but affect services after the migration.

| No.    | Parameter/Option | Community              |                   | Function | Remarks |
|--------|------------------|------------------------|-------------------|----------|---------|
|        |                  | Original Default Value | New Default Value |          |         |
| Server |                  |                        |                   |          |         |

|   |                                  |                   |                    |                                                                         |                                              |
|---|----------------------------------|-------------------|--------------------|-------------------------------------------------------------------------|----------------------------------------------|
| 1 | character_set_server             | latin1            | utf8mb4            | -                                                                       | Be consistent with the origin default value. |
| 2 | collation_server                 | latin1_swedish_ci | utf8mb4_0900_ai_ci | -                                                                       | Be consistent with the origin default value. |
| 3 | explicit_defaults_for_timestamp  | OFF               | ON                 | Specifies whether to update the timestamp column when a row is updated. | Be consistent with the origin default value. |
| 4 | optimizer_trace_max_mem_size     | 16KB              | 1MB                | -                                                                       | Be consistent with the origin default value. |
| 5 | validate_password_check_username | OFF               | ON                 | -                                                                       | Be consistent with the origin default value. |

|        |                         |                                                                                             |                                                                          |                                                                                                                                                                                                 |                                                             |
|--------|-------------------------|---------------------------------------------------------------------------------------------|--------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------|
| 6      | back_log                | -1<br>(autosize)<br>changed<br>from :<br>back_log =<br>50 +<br>(max_con<br>nections /<br>5) | -1<br>(autosize)<br>changed<br>to :<br>back_log =<br>max_conn<br>ections | Specifies<br>the<br>number of<br>requests<br>that can<br>be stored<br>in the<br>stack in a<br>short<br>period<br>before the<br>MySQL<br>database<br>stops<br>respon<br>ding to new<br>requests. | Be<br>consistent<br>with the<br>origin<br>default<br>value. |
| 7      | max_allow<br>ed_packet  | 4194304<br>(4MB)                                                                            | 67108864<br>(64MB)                                                       | Limits the<br>size of<br>data<br>packets<br>received<br>by the<br>server                                                                                                                        | Use the<br>default<br>value.                                |
| 8      | max_error<br>_count     | 64                                                                                          | 1024                                                                     | Controls<br>the<br>number of<br>alarms to<br>be<br>displayed.                                                                                                                                   | Be<br>consistent<br>with the<br>origin<br>default<br>value. |
| 9      | event_sch<br>eduler     | OFF                                                                                         | ON                                                                       | -                                                                                                                                                                                               | Be<br>consistent<br>with the<br>origin<br>default<br>value. |
| 10     | table_ope<br>n_cache    | 2000                                                                                        | 4000                                                                     | -                                                                                                                                                                                               | Be<br>consistent<br>with the<br>origin<br>default<br>value. |
| 11     | log_error_<br>verbosity | 3 (Notes)                                                                                   | 2<br>(Warning)                                                           | -                                                                                                                                                                                               | Use the<br>default<br>value.                                |
| INNODB |                         |                                                                                             |                                                                          |                                                                                                                                                                                                 |                                                             |



|   |                                |                 |                                    |                                                                                                                  |                                                   |
|---|--------------------------------|-----------------|------------------------------------|------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|
| 1 | innodb_undo_tablespaces        | 0               | 2                                  | -                                                                                                                | Use the default value.                            |
| 2 | innodb_undo_log_truncate       | OFF             | ON                                 | -                                                                                                                | Use the default value.                            |
| 3 | innodb_flush_method            | NULL            | fsync (Unix), unbuffered (Windows) | Controls the enabling and writing modes of InnoDB data files and redo logs.                                      | Use the default value<br><b>O_DIRECT</b> for SQL. |
| 4 | innodb_autoinc_lock_mode       | 1 (consecutive) | 2 (interleaved)                    | Controls the behavior of related locks when data is inserted into a table with the <b>auto_increment</b> column. | Be consistent with the origin default value.      |
| 5 | innodb_flush_neighbors         | 1 (enable)      | 0 (disable)                        | Checks whether other dirty pages in the same range are refreshed when refreshing the page from the buffer pool.  | Be consistent with the origin default value.      |
| 6 | innodb_max_dirty_pages_pct_lwm | 0 (%)           | 10 (%)                             | Affects the InnoDB dirty page refreshing operation.                                                              | Use the default value.                            |

|                     |                                  |                        |                       |                                                     |                                                    |
|---------------------|----------------------------------|------------------------|-----------------------|-----------------------------------------------------|----------------------------------------------------|
| 7                   | innodb_max_dirty_pages_pct       | 75 (%)                 | 90 (%)                | Affects the InnoDB dirty page refreshing operation. | Use the default value.                             |
| PERFORM ANCE SCHEMA | Enabled globally.                | -                      | -                     | -                                                   | Be consistent with the origin default value.       |
| REPLICATION         |                                  |                        |                       |                                                     |                                                    |
| 1                   | log_bin                          | OFF                    | ON                    | -                                                   | Enabled by default                                 |
| 2                   | server_id                        | 0                      | 1                     | -                                                   | If the value is <b>0</b> , change it to <b>1</b> . |
| 3                   | log-slave-updates                | OFF                    | ON                    | -                                                   | Enabled by default.                                |
| 4                   | expire_log_days                  | 0                      | 30                    | -                                                   | Use the default value.                             |
| 5                   | master-info-repository           | FILE                   | TABLE                 | -                                                   | Use the default value <b>TABLE</b> .               |
| 6                   | relay-log-info-repository        | FILE                   | TABLE                 | -                                                   | Use the default value <b>TABLE</b> .               |
| 7                   | transaction-write-set-extraction | OFF                    | XXHASH64              | -                                                   | Use the default value.                             |
| 8                   | slave_rows_search_algorithms     | INDEX_SCAN, TABLE_SCAN | INDEX_SCAN, HASH_SCAN | -                                                   | Use the default value.                             |

- Remove system variables.

The analysis of MySQL 5.7 and 8.0 Community Edition shows that removing system variables does not affect migration.

| <b>System variables</b>                |
|----------------------------------------|
| innodb_locks_unsafe_for_binlog         |
| log_built_in_as_identified_by_password |
| old_passwords                          |
| query_cache_limit                      |
| query_cache_min_res_unit               |
| query_cache_size                       |
| query_cache_type                       |
| query_cache_wlock_invalidate           |
| ndb_cache_check_time                   |
| ignore_db_dirs                         |
| tx_isolation                           |
| tx_read_only                           |
| sync_frm                               |
| secure_auth                            |
| multi_range_count                      |
| log_error_verbosity                    |
| sql_log_bin                            |
| metadata_locks_cache_size              |
| metadata_locks_hash_instances          |
| date_format                            |
| datetime_format                        |
| time_format                            |
| max_tmp_tables                         |
| ignore_built_in_innodb                 |
| innodb_support_xa                      |
| innodb_undo_logs                       |
| innodb_undo_tablespaces                |
| internal_tmp_disk_storage_engine       |

## 5.4.8 What Can I Do When OOM Occurs During the Migration of MongoDB Databases?

### Scenarios

Out of memory (OOM) occurs during the migration of MongoDB databases, causing migration failures.

### Possible Cause

The possible causes are as follows:

- If the mongod service of the source database is deployed on a single server, OOM occurs when the migration process consumes large amounts of memory through operations such as creating indexes and sorting queries.
- If the mongod service is deployed on a server with other services and the **cacheSizeGB** value is not specified, OOM occurs when all available memory has been allocated to other services, so the WiredTiger engine does not have sufficient memory.

#### NOTE

By default, the memory used by the WiredTiger engine of mongod is calculated as follows: (Memory in GB-1) x 50% for version 3.2 or (Memory in GB-1) x 50% for version 3.4 and later.

### Solution

- If the mongod service is deployed on a single server, do not perform any operations that consume large amounts of memory during the migration.
- If the mongod service and other services are deployed on the same server, set the value of **cacheSizeGB** to the half of the minimal idle memory to ensure that memory used in peak hours will not be allocated to WiredTiger excessively.

## 5.4.9 How Do I Disable the Balancer?

Before using the DRS service to migrate collections between sharded clusters, you must disable the balancer of the collections to be migrated.

#### NOTE

- After the migration is complete, enable the balancer. The balancer is disabled during the migration, generating different numbers of chunks on each shard of the source database. After the balancer is enabled, chunks will be distributed between shards in the cluster, which may affect the performance of the source database.

### Procedure

**Step 1** Log in to a database through mongo shell.

**Step 2** Run the following command in the command window of the mongos node to switch to the config database:

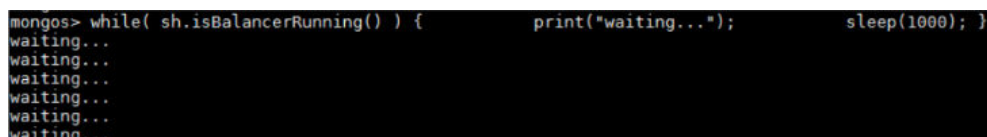
**use config**

**Step 3** Run the following commands to check whether the balancer can be disabled:

```
while(sh.isBalancerRunning()) {
 print("waiting...");
 sleep(1000);
}
```

- If the command output is **waiting**, the balancer is migrating chunks. In this case, do not disable the balancer. Otherwise, data inconsistency may occur.

**Figure 5-4** Viewing the command output



```
mongos> while(sh.isBalancerRunning()) { print("waiting..."); sleep(1000); }
waiting...
waiting...
waiting...
waiting...
waiting...
waiting...
```

- If no command output is displayed, the balancer is not migrating any chunks. In this case, you can disable the balancer:

**Step 4** Disable the balancer.

- If you migrate the entire DB instance, run the following command to disable the balancer.

```
sh.stopBalancer()
```

- If you need to disable the balancer of the sharded collections to be migrated, run the following command:

```
sh.disableBalancing("database.collection")
```

**database.collection** indicates the namespace of the collection to be migrated.

----End

## 5.4.10 How Do I Export and Import Events and Triggers in Batches?

During the MySQL to MySQL migration, if the migration log indicates that the migration of events and triggers fails after the migration task is complete, you can manually migrate the events and triggers.

This section describes how to export and import events and triggers in batches.

**Step 1** Export triggers from the source database in batches.

1. Run the following statement in the source database to obtain values of **TRIGGER\_SCHEMA** and **TRIGGER\_NAME**:

```
SELECT TRIGGER_SCHEMA,TRIGGER_NAME FROM INFORMATION_SCHEMA.TRIGGERS
WHERE TRIGGER_SCHEMA in ('DB1','DB2','DB3') order by TRIGGER_NAME;
```

In the preceding statements, **DB1**, **DB2**, and **DB3** indicate the databases to be migrated to the destination database.

2. Run the following statement in the source database to obtain the statement for creating a trigger from the source database from the **SQL Original Statement** field:

```
SHOW CREATE TRIGGER TRIGGER_SCHEMA.TRIGGER_NAME \G;
```

In the preceding statement, replace **TRIGGER\_SCHEMA** and **TRIGGER\_NAME** with the values obtained in [Step 1.1](#).

**Step 2** Export events from the source database in batches.

1. Run the following statement in the source database to obtain values of **EVENT\_SCHEMA** and **EVENT\_NAME**:  

```
SELECT EVENT_SCHEMA,EVENT_NAME FROM INFORMATION_SCHEMA.EVENTS WHERE
EVENT_SCHEMA in ('DB1','DB2','DB3') order by EVENT_NAME;
```

In the preceding statements, **DB1**, **DB2**, and **DB3** indicate the databases to be migrated to the destination database.
2. Run the following statement in the source database to obtain the statement for creating an event from the source database from the **SQL Original Statement** field:  

```
SHOW CREATE EVENT EVENT_SCHEMA.EVENT_NAME \G;
```

In the preceding statement, replace **EVENT\_SCHEMA** and **EVENT\_NAME** with the values obtained in [Step 2.1](#).

**Step 3** Import triggers and events.

Execute the statements for creating triggers and events exported from the source database in the destination database.

----End

## 5.4.11 How Can I Migrate Databases or Tables Whose Names Contain Uppercase Letters?

### Scenarios

When the value of source database parameter **lower\_case\_table\_names** is set to **1**, the databases or tables whose names contain uppercase letters cannot be migrated.

### Possible Cause

When the value of **lower\_case\_table\_names** in the source database is **1**, the MySQL engine converts the database name or table name into lowercase letters. In this case, the database or table may not be found, resulting in query failure. Simply, if the value of **lower\_case\_table\_names** is **1**, the database or table containing uppercase letters may be inaccessible.

### Solutions

Two solutions are provided as follows:

#### Solution 1

Change the value of **lower\_case\_table\_names** in the source database to **0** (case-sensitive) and ensure that the value of this parameter in the source database is the same as that in the destination database.

#### Solution 2

If the value of **lower\_case\_table\_names** cannot be changed permanently, change the value to **0**, and then perform the following operations:

- For a table, you can use the following statement to convert the table name to lowercase:

```
alter table `BigTab` rename to `bigtab`
```

- For a database, you need to export the database data, change the database name from uppercase to lowercase, and then import the data.

#### CAUTION

After changing the database name or table name, you need to maintain the permission consistency without affecting application access.

## Method 3

Do not migrate the databases or tables that contain uppercase letters.

## 5.4.12 How Do I Delete Orphaned Documents in MongoDB Sharded Clusters?

### What Is Orphaned Document?

In a sharded cluster, orphaned documents are those documents on a shard that also exist in chunks on other shards as a result of failed migrations or incomplete migration cleanup due to abnormal shutdown.

### Migration Impact

During cluster migration, DRS extracts full data from shards. Normal documents and orphaned documents are on different shards and DRS will migrate them all. If the conflict policy of DRS for MongoDB migration is **Ignore**, documents that are first migrated to the destination are stored, resulting in data inconsistency.

### Procedure

**Step 1** Modify the **cleanupOrphaned.js** script file and replace **test** with the database name of the orphaned document to be cleared.

**Step 2** Run the following command to clear the orphaned documents of all collections in the specified database on the shard node:

```
mongo --host ShardIP --port Primaryport --authenticationDatabase database -u username -p password
cleanupOrphaned.js
```

#### NOTE

- **ShardIP**: indicates the IP address of the shard node.
- **Primaryport**: indicates the service port of the primary shard node.
- **database**: indicates the database name.
- **username**: indicates the username for logging in to the database.
- **password**: indicates the password for logging in to the database.

 NOTE

If you have multiple databases, repeat [Step 1](#) and [Step 2](#) to clean up orphaned documents in each database on each shard node.

----End

## 5.5 Backup Migration

### 5.5.1 What Should I Do If the Last Backup File Is Incorrectly Selected in the Backup Migration Scenario?

During the backup migration, if **Last Backup File** is selected by mistake, perform either of the following operations:

- If you select **Yes** by mistake, the database receives a signal that the restore is complete, and then sets the database to available, making incremental backup migration impossible. In this case, you can only delete the backup database and perform full and incremental backup restoration again.
- SQL Server does not have the last backup file in a strict sense. If you select **No** by mistake, you can perform an incremental backup (even if no data is changed). During the incremental backup, select **Yes** to complete the migration. The related database becomes available.

### 5.5.2 Manual Configuration

#### Scenarios

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

#### Login Account

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

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

```
SELECT 'IF (SUSER_ID('+QUOTENAME(SP.name,'''')+') IS NULL) BEGIN CREATE LOGIN '+
+QUOTENAME(SP.name)+
CASE
WHEN SP.type_desc = 'SQL_LOGIN' THEN ' WITH PASSWORD = '
+CONVERT(NVARCHAR(MAX),SL.password_hash,1)+' HASHED,SID='
+CONVERT(NVARCHAR(MAX),SP.SID,1)+' ,CHECK_EXPIRATION = '
```



```
+ CASE WHEN SL.is_expiration_checked = 1 THEN 'ON' ELSE 'OFF' END +', CHECK_POLICY = '
+CASE WHEN SL.is_policy_checked = 1 THEN 'ON,' ELSE 'OFF,' END
ELSE ' FROM WINDOWS WITH'
END
+' DEFAULT_DATABASE=[' +SP.default_database_name+ '], DEFAULT_LANGUAGE=['
+SP.default_language_name+ '] END;' as CreateLogin
FROM sys.server_principals AS SP LEFT JOIN sys.sql_logins AS SL
ON SP.principal_id = SL.principal_id
WHERE SP.type = 'S'
AND SP.name NOT LIKE '##%##'
AND SP.name NOT LIKE 'NT AUTHORITY%'
AND SP.name NOT LIKE 'NT SERVICE%'
AND SP.name NOT IN ('rdsadmin','rdsbackup','rdsuser','rdsmirror','public')
```

**Step 2** Execute the script in [Step 1](#):

**Figure 5-5** Obtaining the script



```
CreateLogin
1 IF (SUSER_ID('sa') IS NULL) BEGIN CREATE LOGIN [sa] WITH PASSWORD = 0x010039BF2EFAD6A3DB4E2ABB941E8ED32E5189A4BE757...
2 IF (SUSER_ID('rdsuser2') IS NULL) BEGIN CREATE LOGIN [rdsuser2] WITH PASSWORD = 0x0100E8BCBC25FC67008D4EE75AD660D1...
3 IF (SUSER_ID('csidbo') IS NULL) BEGIN CREATE LOGIN [csidbo] WITH PASSWORD = 0x0100A508789C15CF688664B162A5EDF4F4D2E...
4 IF (SUSER_ID('TestLogin7') IS NULL) BEGIN CREATE LOGIN [TestLogin7] WITH PASSWORD = 0x010073DA9A79E6677E8AF7077EF67...
5 IF (SUSER_ID('rdsuser3') IS NULL) BEGIN CREATE LOGIN [rdsuser3] WITH PASSWORD = 0x01009448FEDBC8E8D5E5E2529384029CA0...
6 IF (SUSER_ID('Test2') IS NULL) BEGIN CREATE LOGIN [Test2] WITH PASSWORD = 0x0100130953CEA8C997D08B66AF65F84E8CAA44...
7 IF (SUSER_ID('Test3') IS NULL) BEGIN CREATE LOGIN [Test3] WITH PASSWORD = 0x0100EE98873948E02595EDCD953B426637281E7...
8 IF (SUSER_ID('Test4') IS NULL) BEGIN CREATE LOGIN [Test4] WITH PASSWORD = 0x01000EE91B9EF087741F16A44E70AA813D0BA8B...
9 IF (SUSER_ID('Test5') IS NULL) BEGIN CREATE LOGIN [Test5] WITH PASSWORD = 0x010056E8F845DF098D2DF9395AF7B7618A20735...
```

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

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

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

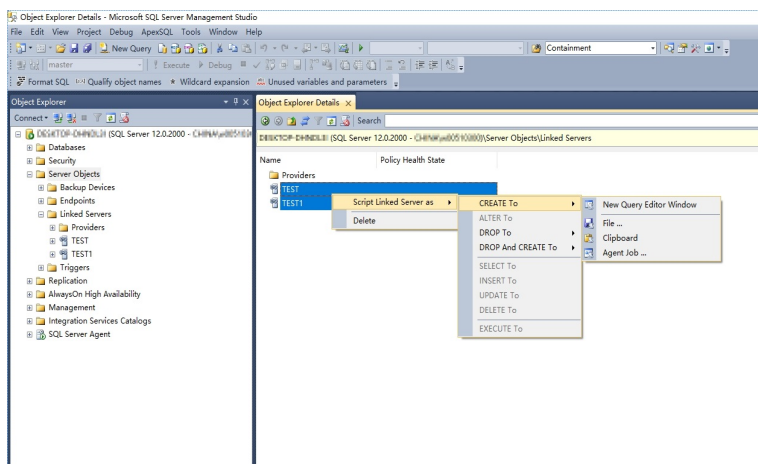
- Step 1** Connect the local DB instance and cloud DB instance through Microsoft SQL Server Management Studio. Choose **Server Objects > Linked Servers** and locate the DBLink of the current DB instance.

**Figure 5-6** 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 5-7** Creating the script



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

```
USE [master]
GO
```

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

 **NOTE**

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


----End

## Agent JOB

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

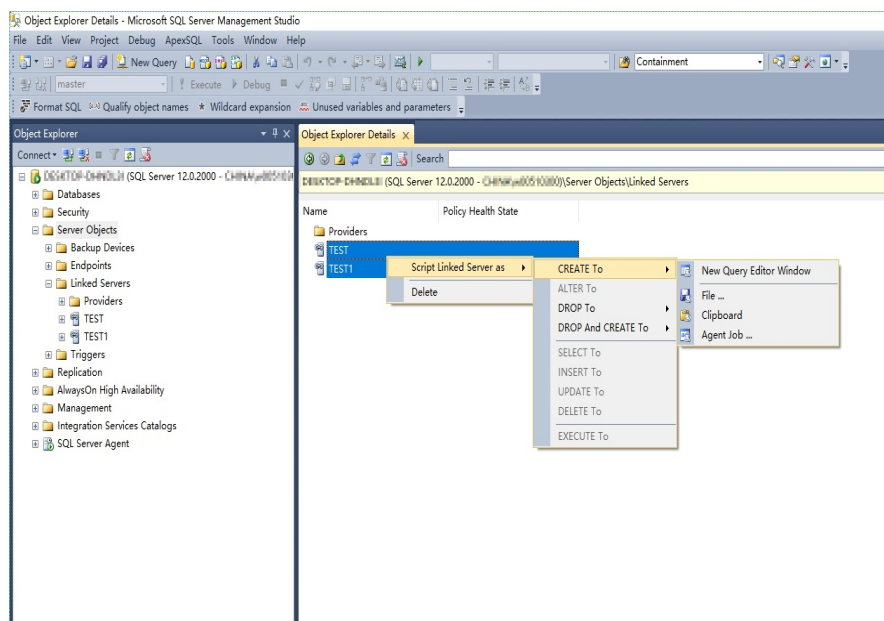
- Step 1** Connect the local DB instance and cloud DB instance through Microsoft SQL Server Management Studio. Choose **SQL Server Agent > Jobs** and locate all the jobs of the current DB instance.

**Figure 5-8** 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 5-9 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 SQL Server DB instance, only the owner of the job can view the job of the DB instance. Therefore, it is recommended that all job owners use the same account to facilitate job management.

----End

## Key Configuration

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

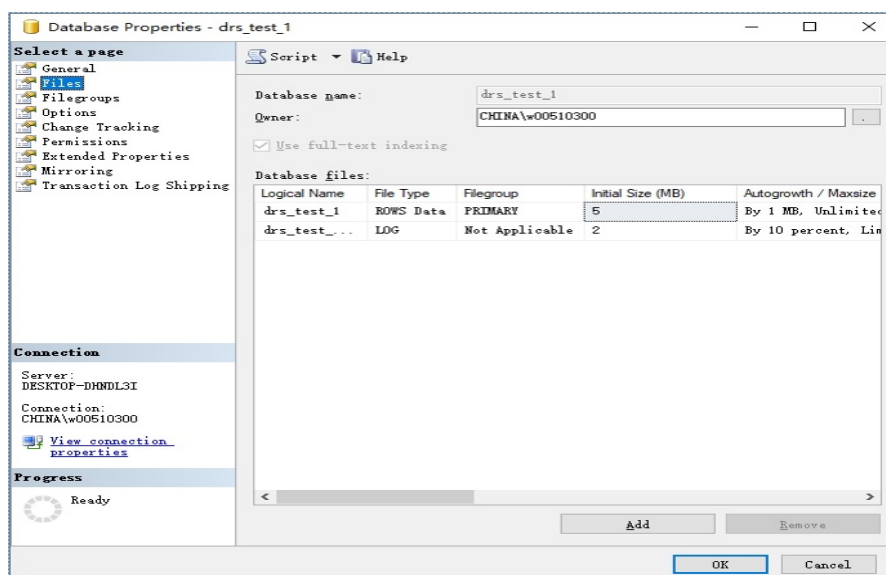
1. tempdb: The file configuration of the temporary database needs to be synchronized.

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

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

```
USE [master]
GO
```

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

**Figure 5-10** Checking temporary files

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


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

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

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


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

Figure 5-11 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 5-12 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 5-13 Checking the database recovery model



## 5.6 Real-Time Synchronization

### 5.6.1 Can DRS Sync Tables of Different Schemas to the Same Schema?

DRS can directly synchronize tables of different schemas to those of the same schema if the tables do not conflict with each other.

### 5.6.2 Can Online DDL Tools Be Used for Real-time Synchronization?

During table-level incremental synchronization from MySQL to MySQL, you can use Online DDL tools to add or delete columns. Pay attention to the following when using Online DDL:

- The DRS synchronization mechanism conflicts with the tool. Do not select **Incremental DDLs** when selecting synchronization objects on the **Set Synchronization Task** page.
- Before using Online DDL tools to add columns, perform the corresponding operations in the destination database and then in the source database.
- When using Online DDL tools to delete columns, perform the corresponding operation in the source database and then in the destination database.

Common online DDL tools:

- pt-online-schema-change
- gh-ost

### 5.6.3 Which MySQL DDL Statements Can Be Directly Executed in GaussDB(DWS)?

During the real-time synchronization from MySQL to GaussDB(DWS), DRS does not convert incremental MySQL DDL statements, but directly executes these DDL statements in GaussDB(DWS). The following DDL statements can be directly executed in GaussDB(DWS):

```
alter table test add column c1 varchar(20);
alter table test drop column c1;
drop table test;
```



```
create table test(id int primary key ,name varchar(20));
create index test on db1.test (name);
truncate table db1.test;
```

The following DDL statements cannot be directly executed:

```
alter table db1.test rename to db1.test_bak;
rename table test to test_bak;
alter table db1.test modify column c1 varchar(40);
alter table db1.test change column c1 c1 varchar(50);
alter table db1.test add index/key index_1(name);
alter table db1.test drop index/key index_1;
```

## 5.6.4 Why Do I Use the SCAN IP Address to Connect to an Oracle RAC Cluster?

If the source Oracle database is an RAC cluster, you are advised to use SCAN IP +SERVICE\_NAMES to create a task because SCAN IP has stronger fault tolerance, better load balancing capability, and faster synchronization.

- If the SCAN IP address is used, ensure that the SCAN IP address can communicate with all virtual IP addresses of the source database. Otherwise, the connection test cannot be passed.
- If SCAN IP is not used, the virtual IP address of a node can be used. If other nodes are abnormal, the synchronization process is not affected.

For details about the SCAN IP address, see the [documents](#) on the Oracle official website.

## 5.6.5 How Do I Check Supplemental Logging of the Source Oracle Database?

In physical standby mode, the Oracle database directly replicates logs from the primary database and does not generate any logs. If the source is an Oracle database, you need to check whether supplemental logging on the primary database meets the requirements to ensure that the task can run properly. The following lists the check and setting methods:

Table level: This setting applies to a specified table.

Database level: This setting applies to the database level.

PK/UI: In addition to the changed columns, the values of the primary key and unique key of each row are recorded.

ALL: Each row of the log records the values of all columns in that row.

### NOTE

DRS incremental synchronization requirements can be met if any of the following checks are passed.

### Table-level PK/UI Supplemental Logging Check (Minimum Requirement)

Check whether supplemental logging of the table-level objects to be synchronized meets the requirements.

**Step 1** Run the following SQL statement in the source database:

```
select * from ALL_LOG_GROUPS where (LOG_GROUP_TYPE='UNIQUE KEY LOGGING' or
LOG_GROUP_TYPE='PRIMARY KEY LOGGING') and OWNER='Schema name in uppercase' and
TABLE_NAME='Table name in uppercase';
```

If the table name corresponds to the records whose **LOG\_GROUP\_TYPE** is **UNIQUE KEY LOGGING** and **PRIMARY KEY LOGGING** in the query result, the DRS incremental synchronization requirements are met.

**Step 2** If the requirements are not met, run the following SQL statement to enable table-level PK/UI logging:

```
alter database add supplemental log data;
alter table Schema_name.Table_name add supplemental log data(primary key,unique) columns;
```

----End

## All Table-Level Supplemental Log Check

Check whether supplemental logging of the table-level objects to be synchronized meets the requirements.

**Step 1** Run the following SQL statement in the source database:

```
select * from ALL_LOG_GROUPS where LOG_GROUP_TYPE='ALL COLUMN LOGGING' and
OWNER='Schema_name in uppercase' and TABLE_NAME='Table name in uppercase';
```

If the table name is recorded in the query result, the DRS incremental synchronization requirements can be met.

**Step 2** If the requirements are not met, run the following SQL statement to enable all column supplemental logging at the table level:

```
alter database add supplemental log data;
alter table Schema_name.Table_name add supplemental log data(all) columns;
```

----End

## Database-level Supplemental Log Check

For the database-level objects to be synchronized, check whether supplemental logging meets the requirements.

**Step 1** Run the following SQL statement in the source database:

```
select SUPPLEMENTAL_LOG_DATA_MIN MIN, SUPPLEMENTAL_LOG_DATA_PK PK,
SUPPLEMENTAL_LOG_DATA_UI UI, SUPPLEMENTAL_LOG_DATA_ALL ALL_LOG from v$database;
```

**Step 2** Either of the following requirements must be met:

- If both **PK** and **UI** are set to **YES**, DRS incremental synchronization requirements can be met.

If the requirements are not met, run the following SQL statement to enable database-level PK/UI supplemental logging:

```
alter database add supplemental log data(primary key, unique) columns;
```

- If **ALL\_LOG** is set to **YES**, DRS incremental synchronization requirements can be met.

If the requirements are not met, run the following SQL statement to enable all column supplemental logging at the database level:

```
alter database add supplemental log data(all) columns;
```

----End

## 5.6.6 Garbled Characters or Synchronization Failure Due to Incompatible Character Sets

If the character set of the source database is incompatible with that of the destination database, some data may include garbled characters, data synchronization may be inconsistent, or data may fail to be written into the destination database. In this case, change the character set of the destination database before synchronization.

## 5.7 Data-Level Comparison

### 5.7.1 Which of the Following Data Types Are Not Supported By Value Comparison?

DRS's data comparison allows you to check whether the data in the source database is the same as that in the destination database.

DRS does not support value comparison for the data types shown here. During value comparison, these data types are automatically ignored.

**Table 5-2** Data types that do not support value comparison

| Source DB Type | Data Type                                                                                                               |
|----------------|-------------------------------------------------------------------------------------------------------------------------|
| MySQL          | TINYBLOB, BLOB, MEDIUMBLOB, LONGBLOB, TINYTEXT, TEXT, MEDIUMTEXT, LONGTEXT                                              |
| GaussDB        | TEXT, CLOB, BLOB, BYTEA, INTERVAL DAY TO SECOND, INTERVAL                                                               |
| Oracle         | BLOB, NCLOB, CLOB, LONG RAW, LONG, INTERVAL DAY TO SECOND, INTERVAL YEAR TO MONTH, UROWID, BFILE, XMLTYPE, SDO_GEOMETRY |
| MongoDB        | _id is of the bindata type.                                                                                             |

DRS does not support value comparison for the following primary key types. During value comparison, the following primary key types are grouped into a specified table that does not support comparison.

**Table 5-3** Primary key type that does not support value comparison.

| Source DB Type | Data Type                                                                         |
|----------------|-----------------------------------------------------------------------------------|
| MySQL          | TINYBLOB, BLOB, MEDIUMBLOB, LONGBLOB, TINYTEXT, TEXT, MEDIUMTEXT, LONGTEXT, FLOAT |

| Source DB Type | Data Type                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GaussDB        | TEXT, CLOB, BLOB, BYTEA, INTERVAL DAY TO SECOND, INTERVAL, REAL, DOUBLE PRECISION, BOOL, TIME, TIMETZ, TIMESTAMP, TIMESTAMPTZ, DATE                                                                                                                                                                                                                                                                            |
| Oracle         | BLOB, NCLOB, CLOB, LONG RAW, LONG, INTERVAL DAY TO SECOND, INTERVAL YEAR TO MONTH, UROWID, BFILE, XMLTYPE, SDO_GEOMETRY, BINARY_FLOAT, BINARY_DOUBLE, FLOAT, RAW, TIMESTAMP, TIMESTAMP WITH TIME ZONE, TIMESTAMP WITH LOCAL TIME ZONE, DATE                                                                                                                                                                    |
| PostgreSQL     | REAL, DOUBLE PRECISION, MONEY, TEXT, BYTEA, TIMESTAMP WITHOUT TIME ZONE, TIMESTAMP WITH TIME ZONE, DATE, TIME WITHOUT TIME ZONE, TIME WITH TIME ZONE, INTERVAL, BOOLEAN, ENUMERATED TYPES, POINT, LINE, LSEG, BOX, PATH, POLYGON, CIRCLE, CIDR, INET, MACADDR, MACADDR8, BIT, BIT VARYING, TSVECTOR, TSQUERY, XML, JSON, ARRAY, COMPOSITE TYPES, INT4RANGE, INT8RANGE, NUMRANGE, TSRANGE, TSTZRANGE, DATERANGE |

## 5.7.2 What Impact Does a DRS Comparison Task Have on Databases?

- Object comparison: System tables of the source and destination databases are queried, occupying about 10 sessions. The database is not affected. However, if there are a large number of objects (for example, hundreds of thousands of tables), the database may be overloaded.
- Row comparison: The number of rows in the source and destination databases is queried, which occupies about 10 sessions. The SELECT COUNT statement does not affect the database. However, if a table contains a large amount of data (hundreds of millions of records), the database will be overloaded and the query results will be returned slowly.
- Value comparison: All data in the source and destination databases is queried, and each field is compared. The query pressure on the database leads to high I/O. The query speed is limited by the I/O and network bandwidth of the source and destination databases. Value comparison occupies one or two CPUs, and about 10 sessions.

## 5.7.3 How Long Does a DRS Comparison Task Take?

- Object comparison: Generally, the comparison results are returned within several minutes based on the query performance of the source database. If the amount of data is large, the comparison may take dozens of minutes.
- Row comparison: The SELECT COUNT method is used. The query speed depends on the database performance.

- Value comparison: If the database workload is not heavy and the network is normal, the comparison speed is about 5 MB/s.

## 5.8 General Operations

### 5.8.1 What Can I Do When Information Overlaps on the DRS Console?

Information often overlaps when you decrease the size of the page. You are advised to set the page scale at 100%.

### 5.8.2 Is the Destination Instance Set to Read-only or Read/Write?

When configuring a migration task, you can set the destination instance to **Read-only** or **Read/Write**.

- **Read-only:** During the migration, the entire destination instance is read-only. After the migration is complete, it restores to the read/write status. This option ensures the integrity and success rate of data migration.
- **Read/Write:** During the migration, the destination instance can be queried or modified. Data being migrated may be modified when operations are performed or applications are connected. It should be noted that background processes can often generate or modify data, which may result in data conflicts, task faults, and upload failures. Do not select this option if you do not fully understand the risks.

Setting the destination instance to read-only can prevent DDL or DML misoperations from being performed on the databases or tables that are being migrated, improving migration integrity and data consistency.

- After a migration task is started, the status of the destination database cannot be changed.
- After all migration tasks in which the destination database status is set to read-only are complete, the destination database can be read and written.

### 5.8.3 How Do I Set Global `binlog_format=ROW` to Take Effect Immediately?

During migration for MySQL databases, the source database binlog must be in the ROW format. Otherwise, the task fails. After **`binlog_format=ROW`** at the global level is set in the source database, all the previous service threads need to be stopped because these threads still connect the binlog in the non-ROW format.

#### Procedure

- Step 1** Log in to the source database using the MySQL official client or other tools.
- Step 2** Run the following command for setting global parameters in the source database.  

```
set global binlog_format = ROW;
```

**Step 3** Run the following command on the source database and check whether the preceding operation is successful:

```
select @@global.binlog_format;
```

**Step 4** You can use either of the following methods to ensure that the modified binlog format of the source database takes effect immediately:

#### Method 1

1. Select a non-service period to disconnect all service connections on the current database.
  - a. Run the following command to query all service threads (excluding all binlog dump threads and current threads) in the current database:

```
show processlist;
```
  - b. Stop all the service threads queried in the previous step.

#### NOTE

- Do not create or start a migration task before the preceding operations are complete. Otherwise, data may be inconsistent.
2. To prevent the binlog format of the source database from becoming invalid due to database restart, add or modify the **binlog\_format** parameter in the startup configuration file (**my.ini** or **my.cnf**) of the source database and save the modification.

```
binlog_format=ROW
```

#### Method 2

1. To prevent the binlog format of the source database from becoming invalid due to database restart, add or modify the **binlog\_format** parameter in the startup configuration file (**my.ini** or **my.cnf**) of the source database and save the modification.

```
binlog_format=ROW
```
2. Ensure that the **binlog\_format** parameter is successfully added or modified. Then, restart the source database at a non-service period.

----End

## 5.8.4 How Do I Set binlog\_row\_image=FULL to Take Effect Immediately?

When migrating MySQL databases, ensure that the **binlog\_row\_image** parameter of the source database is set to **FULL**. Otherwise, the migration task will fail. After **binlog\_row\_image** is set to **FULL** in the source database, the setting takes effect only for new sessions. To close old sessions, restart the source database and reset the task during a non-service period.

### Setting binlog\_row\_image to FULL

- If the source is an RDS instance on the cloud, change **binlog\_row\_image** to **FULL** on the RDS console, and then restart the source database and reset the task.
- If the source database is an on-premises database, perform the following steps:

- a. Log in to the server where the MySQL source database is located.
- b. Manually change the value of **binlog\_row\_image** in the **my.cnf** configuration file to **FULL** and save the file.  
`binlog_row_image=full`
- c. To close old sessions, restart the source database and reset the task during a non-service period.

## 5.8.5 How Do I Change the Destination Database Password to Meet the Password Policy?

### Scenarios

When you set the password for the migration account in the destination database, you need to set the password based on the password strength requirements of the destination database.

### Procedure

The following operations apply to the scenario where the target database is an RDS instance.

- Step 1** Log in to the RDS console.
- Step 2** Locate the target DB instance.
- Step 3** Click the DB instance name.
- Step 4** On the **Basic Information** page, click the **Parameters** tab.
- Step 5** Enter the keyword **password** in the search box in the upper right corner of the page and press **Enter** to view the search result.
- Step 6** In the search result in **Step 5**, change the values of the parameters listed in **Table 5-4** based on the password strength requirements. Ensure that the parameter values are within the password complexity range.

**Table 5-4** Password description

| Parameter                          | Allowed Value   | Description                                                                                                                              |
|------------------------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------|
| validate_password_length           | 0-2,147,483,647 | Specifies the minimum password length verified by the validate_password plugin.                                                          |
| validate_password_mixed_case_count | 0-2,147,483,647 | Specifies the minimum number of lowercase and uppercase letters in a password when the password policy level is <b>MEDIUM</b> or higher. |
| validate_password_number_count     | 0-2,147,483,647 | Specifies the minimum number of digits in a password when the password policy level is <b>MEDIUM</b> or higher.                          |

| Parameter                            | Allowed Value       | Description                                                                                                                          |
|--------------------------------------|---------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| validate_password_policy             | LOW, MEDIUM, STRONG | Specifies the password policy executed by the validate_password plugin.                                                              |
| validate_password_special_char_count | 0-2,147,483,647     | Specifies the minimum number of non-alphanumeric characters in a password when the password policy level is <b>MEDIUM</b> or higher. |

**Step 7** After the parameter values are modified, save the modification.

**Step 8** Back to the **Select Migration Type** page and perform the next step.

----End

## 5.8.6 How Do I Configure the Shard Key for a MongoDB Sharded Cluster?

MongoDB shards data at the collection level, distributing the collection data using shard keys.

You choose the shard key when sharding a collection. Each record contains a shard key, and the shard key is either an indexed field or indexed compound fields. MongoDB database distributes data in different chunks according to the shard key, and distributes chunks evenly among the shards. To divide data chunks by shard key, MongoDB database uses two sharding methods: range-based sharding and hashed sharding.

**Table 5-5** Shard key classification

| Shard Key Type       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                  | Application Scenario                                                                                                                 |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| Range-based sharding | <p>Ranged-based sharding involves dividing data into contiguous ranges determined by the shard key values. Range-based sharding is the default sharding methodology if no other options are specified.</p> <p>This allows for efficient queries where reads target documents within a contiguous range. The distribution route determines which data chunk stores the data required and forwards the request to the corresponding shard.</p> | It is recommended when the shard key has high cardinality with low frequency, and the shard key value does not change monotonically. |



| Shard Key Type  | Description                                                                                                                                                                                                                                                                                                              | Application Scenario                                                                                                                                             |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hashed sharding | Hashed sharding uses a hashed index to partition data across your shared cluster and to create chunks.<br>Hashed sharding provides more even data distribution across the sharded cluster. Hash values enable data to be randomly distributed in each chunk, and therefore are randomly distributed in different shards. | If the shard key values that have a high cardinality or change monotonically, or there are large number of different values, hashed sharding is an ideal option. |

Once you shard a collection, the shard key and the shard key values are immutable. If you need to modify the shard key of a document, you must delete the document. Then modify the shard key and insert the document again.

 NOTE

The shard key does not support array indexes, text indexes, geographical indexes, and spatial indexes.

## Range-based Sharding

**Step 1** Run the following command to enable database sharding:

```
sh.enableSharding(database)
```

 NOTE

*database* indicates the database for which the sharded collection is enabled.

**Step 2** Configure the collection's shard key.

```
sh.shardCollection(namespace, key)
```

 NOTE

- *namespace* consists of a string <database>.<collections> specifying the full namespace of the target collection.
- *key* indicates the index for the shard key.
- If the collection is empty, skip this step because the index on the shard key can be created automatically.

```
sh.shardCollection()
```

- If the collection is not empty, create an index key. Then, run the following command to set the shard key:

```
sh.shardCollection()
```

----End

## Hashed Sharding

**Step 1** Run the following command to enable database sharding:

```
sh.enableSharding(database)
```

 NOTE

*database* indicates the database for which the sharded collection is enabled.

**Step 2** Set hashed shard keys.

```
sh.shardCollection("<database>.<collection>", { <shard key> : "hashed" }, false, {numInitialChunks: Number of preconfigured chunks})
```

The value of **numInitialChunks** is calculated as follows: `db.collection.stats().size / 10*1024*1024*1024`.

If the collection contains data, run the following command to create a hashed index for the hashed key:

```
db.collection.createIndex()
```

Run the following command to create a hashed shard key:

```
sh.shardCollection()
```

----End

### 5.8.7 Does Bandwidth Expansion Affect the Running DRS Tasks?

When the cloud connection bandwidth is expanded, the bandwidth link needs to be re-established and the network is disconnected. Whether the network disconnection affects DRS tasks depends on the network disconnection duration and whether the source database IP address changes. For example, for the MySQL DB engine, if the network is disconnected for one day and the binlog of the source database is cleared within this day (the binlog clearing policy of MySQL is configured by the user), the task cannot be resumed. In this scenario, you need to reset the task. If the network is interrupted for a short period of time and the IP address of the source database in the VPN remains unchanged after the bandwidth link is changed, the system can continue to resume the task.

### 5.8.8 Why Data in MariaDB and SysDB Cannot Be Migrated?

In some MariaDB versions, the SysDB database is used as a system database (similar to the sys database of MySQL 5.7). Therefore, DRS considers the SysDB database as the system database of all MariaDB databases by default (similar to the MySQL, information\_schema, and performance\_schema databases). If the SysDB is a service database, you can apply for a service ticket.

### 5.8.9 Constraints and Operation Suggestions on Many-to-One Scenario

DRS supports many-to-one scenarios during migration of different types of instances and tables to suit your service requirements.

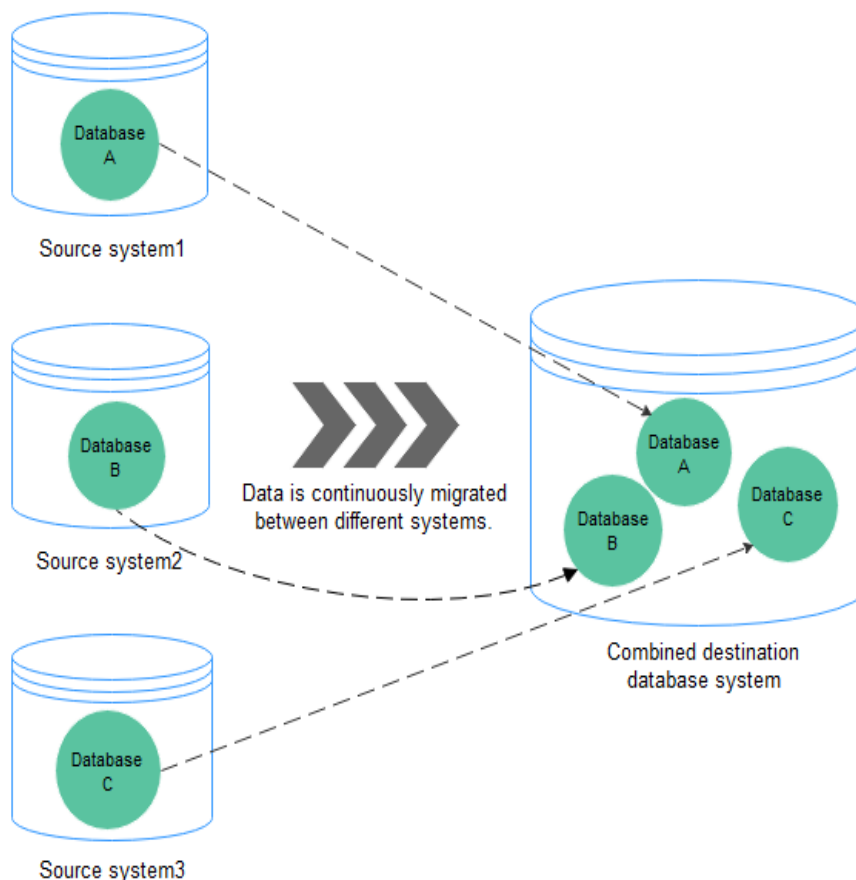
## Operation Suggestions

- To ensure that there is sufficient space during task creation, you are advised to calculate the total data volume of the source database and plan how to allocate the disk space of the destination instance. The remaining disk space must be greater than the total data volume of the source database. For example, if the data volume of source system1 is 1 GB, the data volume of source system2 is 3 GB, and the data volume of source system3 is 6 GB, the remaining disk space of the destination instance must be greater than 10 GB.
- To improve the performance of the destination MySQL database, you are advised to use the **Save Change** function to configure common parameters (except **max\_connections**). For performance parameters, you need to manually change the parameter values based on the specifications of the destination database.
- When you create a many-to-one synchronization task, the task created later may block the task created earlier. This is because each synchronization task involves index creation. When an index is created, a schema lock may occur on the destination database, which blocks the synchronization of other tables in the schema. As a result, the previously created tasks cannot be synchronized. To avoid this problem, you are advised to set **Start Time** to **Start at a specified time** to start a task during off-peak hours.
- For many-to-one synchronization tasks that involve the synchronization of the same table, DDL operations cannot be performed on source databases. Otherwise, all synchronization tasks fail.

## Many-to-One Data Migration

Data migration aims to migrate the entire database. Multiple databases can be migrated at the instance level. Databases with the same name in the source system cannot be migrated and database name mapping is not supported.


**Figure 5-14** Many-to-one data migration



## Flow Chart

When creating a task, ensure that the second task is created after the first task has entered the full migration state.

**Figure 5-15** Flow chart



### 5.8.10 Where Can I View DRS Operation Logs?

You can view DRS operation logs on the Cloud Trace Service (CTS) console.

Click the username in the upper right corner and select **Operation Log** from the drop-down list.

## 5.8.11 Can a Completed Task Be Restarted?

No. DRS cannot restart a completed task.

## 5.8.12 What Are the Differences Between Resetting a Task and Recreating a Task?

You can reset a task when the task is suspended or fails. Resetting a task does not clear the destination database. You can determine whether to clear the destination database based on your requirements. After the task is reset, a full synchronization is performed again. You do not need to configure the task again.

## 5.8.13 What Do I Do After Changing the Password of the Source or Destination Database?

A DRS task may fail due to the change of the password of the source or destination database. To continue the task, rectify the information and then retry the task on the DRS console.

### Procedure

- Step 1** Select a task from the task list and click the task name.
- Step 2** On the **Basic Information** tab, click **Modify Connection Details** in the **Connection Information** area.
- Step 3** In the displayed dialog box, change the passwords of the source and destination databases and click **OK**.

----End

## 5.9 Delay

### 5.9.1 Why Does the Delay of DR Tasks Increase?

#### Causes for an Increase in RTO

Recovery Time Objective (RTO) is duration of time within which transactions on the DRS instance are transmitted and replayed to the destination database during incremental synchronization. If the RTO value is large, transactions to be replayed on the DRS instance are stacked. The possible causes are as follows:

1. After a DR task is initialized, the incremental data generated from the time when the DR task is started to the current time needs to be replayed.
2. Batch operations are performed on service database tables that do not have primary keys. The DR instance is synchronizing tables that do not have primary keys and have a large amount of changed data. To ensure data consistency in tables without primary keys, all operations are recorded. As a result, the operation execution efficiency is lower than that in tables with primary keys. In addition, if the destination table has no index, the data update efficiency is lower.

3. If the DDL operation is performed on the service database, the DR instance can replay data only after the execution of the DDL operation is complete.
4. Frequently executed operations are performed on hot tables in the service database. The DR instance combines the transactions of the hot table and then replays the transactions, reducing frequent operations on the destination database.
5. The access to the DR database is abnormal. As a result, the incremental data cannot be replayed.

## Handling Suggestion

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

**Step 2** On the **Basic Information** page, click the **Disaster Recovery Monitoring** tab to view the changes of RTO.

- If RTO decreases gradually or increases only in a short time, no action is required.
- If RTO keeps increasing, run the following statement in the DR database to check whether there are SQL statements that take a long time to execute or DDL statements that are being executed:  

```
show processlist
```
- If the DR database is abnormal, contact database O&M engineers.

----End

## Causes for an Increase in RPO

Recovery Point Objective (RPO) refers to the time that passes from when a transaction in the service database is submitted and the time when the transaction is synchronized to the DRS instance during the incremental synchronization. If RPO is large, the latest changes made to the service database data have not been extracted to the DR instance. The possible causes are as follows:

1. The network between the service database and the DRS DR instance is unstable. Reading changes in logs from the service database is slow.
2. The service database cannot be accessed. As a result, incremental data cannot be extracted.

## Handling Suggestion

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

**Step 2** On the **Basic Information** page, click the **Disaster Recovery Monitoring** tab to view the changes of RPO.

- If RPO decreases gradually or increases only in a short time, no action is required.
- If the service database is abnormal, contact database O&M engineers.

----End

## 5.9.2 Why Is the Delay High In MongoDB Replication Scenarios?

### Involved Scenarios

- Migration from MongoDB to DDS
- Migration from DDS to MongoDB

### Possible Causes

To ensure the performance of migration, synchronization, or disaster recovery, DRS performs concurrent replay at the collection level in the incremental phase. In the following special cases, DRS supports only single-thread write and does not support concurrent replay:

- The collection index contains a unique key.
- The value of **capped** of the collection attribute is **true**.

If the delay increases, check whether the problem is caused by the preceding reasons.

# 6 Troubleshooting

## 6.1 Solutions to Failed Check Items

### 6.1.1 Disk Space

#### 6.1.1.1 Checking Whether the Destination Database Has Sufficient Storage Space

##### MySQL Migration

**Table 6-1** Checking whether the destination database has sufficient storage space

|                                              |                                                                                                                                                                       |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database has sufficient storage space                                                                                                         |
| <b>Description</b>                           | Check whether the destination database has sufficient storage space. If storage space is insufficient, the migration will fail.                                       |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected. |
|                                              | Failure cause: Insufficient user permissions<br>Handling suggestion: Check whether the database user permissions meet the migration requirements.                     |



|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The destination database does not have sufficient storage space (at least 2.5 times the size of the source database).</p> <p>Handling suggestion: Scale up or clean up the destination database storage space. If you clean up the storage space, you will obtain more space within 2 to 3 minutes.</p> <p><b>NOTE</b><br/>It is recommended that the size of the destination database disk be set to the smaller value of the following two values:</p> <ol style="list-style-type: none"> <li>1. 2.5 times the size of the data to be migrated in the source database.</li> <li>2. The size of the data to be migrated in the source database plus 200 GB.</li> </ol> |
|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## PostgreSQL Synchronization

**Table 6-2** Checking whether the destination database has sufficient storage space

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database has sufficient storage space                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                           | Check whether the destination database has sufficient storage space. If storage space is insufficient, the synchronization will fail.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p> <p>Failure cause: The destination database does not have sufficient storage space (at least 1.5 times the size of the source database).</p> <p>Handling suggestion: Scale up or clean up the destination database storage space. If you clean up the storage space, you will obtain more space within 2 to 3 minutes.</p> <p><b>NOTE</b><br/>It is recommended that the size of the destination database disk be set to the smaller value of the following two values:</p> <ol style="list-style-type: none"> <li>1. 1.5 times the size of the data to be migrated in the source database.</li> <li>2. The size of the data to be migrated in the source database plus 200 GB.</li> </ol> |

## MongoDB Migration

**Table 6-3** Checking whether the destination database has sufficient storage space

|                    |                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the destination database has sufficient storage space                                                                   |
| <b>Description</b> | Check whether the destination database has sufficient storage space. If storage space is insufficient, the migration will fail. |

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|                                              | Failure cause: The destination database does not have sufficient storage space (at least 1.5 times the size of the source database).<br>Handling suggestion: Scale up or clean up the destination database storage space. If you clean up the storage space, you will obtain more space within 2 to 3 minutes.<br><b>NOTE</b><br>It is recommended that the size of the destination database disk be set to the smaller value of the following two values:<br><ol style="list-style-type: none"><li>1.5 times the size of the data to be migrated in the source database.</li><li>The size of the data to be migrated in the source database plus 200 GB.</li></ol> |

### 6.1.1.2 Checking Whether the Destination Server Has Sufficient Storage Space

**Table 6-4** Checking whether the destination server has sufficient storage space

|                                              |                                                                                                                                                                                            |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination server has sufficient storage space                                                                                                                                |
| <b>Description</b>                           | If the destination server's storage space is insufficient, the migration will fail.                                                                                                        |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The amount of data in the source database is greater than the remaining storage space of the destination server.<br>Handling suggestion: Modify the synchronization object. |

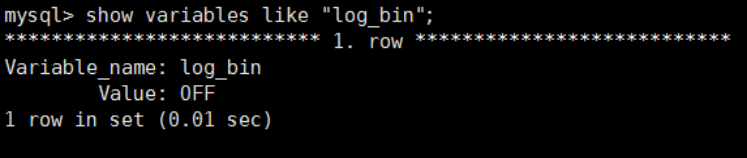
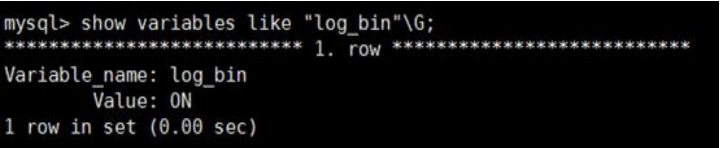
## 6.1.2 Database Parameters

### 6.1.2.1 Checking Whether the Source Database Binlog Is Enabled

#### MySQL Migration

**Table 6-5** Checking whether the source database binlog is enabled

|                   |                                               |
|-------------------|-----------------------------------------------|
| <b>Check Item</b> | Whether the source database binlog is enabled |
|-------------------|-----------------------------------------------|

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>                           | Check whether binlog is enabled for the source database.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p> <p><b>NOTE</b><br/>For details about the required MySQL permissions and authorized operations, see <a href="#">Which MySQL Permissions Are Required for DRS?</a></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|                                              | <p>Failure cause: The binlog function is disabled on the source database.</p> <p>Handling suggestion:</p> <ul style="list-style-type: none"> <li>• If the source is an on-premises database, perform the following operations to enable binlog.             <ol style="list-style-type: none"> <li>1. Run the following command to check whether binlog is enabled:<br/> <pre>show variables like "log_bin";</pre>  </li> <li>2. If binlog is disabled, add <b>log-bin = mysql-bin</b> followed by <b>[mysqld]</b> in the MySQL configuration file <b>my.cnf</b> or <b>my.ini</b>.<br/> <pre>[mysqld] log-bin = mysql-bin</pre> </li> <li>3. Restart the database.<br/>  </li> </ol> </li> <li>• If the source is an RDS instance, enable binlog by referring to the RDS official documentation.</li> </ul> |

### 6.1.2.2 Checking Whether the Source Database Binlog Is Row-Based

#### MySQL Migration

**Table 6-6** Checking whether the source database binlog is row-based

|                   |                                                 |
|-------------------|-------------------------------------------------|
| <b>Check Item</b> | Whether the source database binlog is row-based |
|-------------------|-------------------------------------------------|

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>                           | Check whether the source database binlog is row-based.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|                                              | <p>Failure cause: The source database binlog is not row-based.</p> <p>Handling suggestion:</p> <ul style="list-style-type: none"> <li>If the source database is an on-premises database, perform the following operations to change the binlog format of the source database:<br/>                     Method 1: You can modify the <b>my.cnf</b> or <b>my.ini</b> configuration file and restart the database.<br/> <pre>binlog_format=row</pre>                     Method 2: Stop all service connections.<br/> <pre>set global binlog_format='ROW'</pre>                     Modify the <b>my.cnf</b> or <b>my.ini</b> configuration file.<br/> <pre>binlog_format=row</pre>                     In the <b>ROW</b> format, the log growth rate increases, which may occupy more disk space.                 </li> </ul> <p><b>NOTE</b><br/>                     The <b>MySQL Global binlog_format</b> parameter does not take effect for connected sessions. For details about how to switch format, see <a href="#">How Do I Set Global binlog_format=ROW to Take Effect Immediately?</a></p> <ul style="list-style-type: none"> <li>If the source database is an RDS DB instance on the cloud, change the <b>binlog_format</b> value of the source database to <b>ROW</b>. After the change, restart the database for the change to take effect.</li> </ul> <p><b>NOTE</b><br/>                     The <b>MySQL Global binlog_format</b> parameter does not take effect for connected sessions. For details about how to switch format, see <a href="#">How Do I Set Global binlog_format=ROW to Take Effect Immediately?</a></p> |

### 6.1.2.3 Checking Whether the Binlog Retention Period Is Set on the Source Database

#### MySQL Migration

**Table 6-7** Checking whether the binlog retention period is set on the source database

|                   |                                                                   |
|-------------------|-------------------------------------------------------------------|
| <b>Check Item</b> | Whether the binlog retention period is set on the source database |
|-------------------|-------------------------------------------------------------------|

|                                              |                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>                           | Checking whether the binlog retention period is set on the source database. You are advised to store the source database binlog for a longer time, if the storage space is sufficient.                                                                                                                                                                                          |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The binlog retention period is not set on the source database.</p> <p>Handling suggestion:</p> <p>Log in to the source database and run the following SQL statement to set the retention period of binlog:<br/> <code>call mysql.rds_set_configuration('binlog retention hours', n);</code></p> <p>The value <b>n</b> indicates an integer from 1 to 168.</p> |

### 6.1.2.4 Checking Whether the Source and Destination Database Character Sets Are Consistent

#### MySQL Migration

**Table 6-8** Checking whether the source and destination database character sets are consistent

|                                              |                                                                                                                                                                                            |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source and destination database character sets are consistent                                                                                                                  |
| <b>Description</b>                           | Check whether the character sets of the servers hosting the source and destination databases are consistent.                                                                               |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p>           |
|                                              | <p>Failure cause: This item cannot be checked because the destination database fails to be connected.</p> <p>Handling suggestion: Check whether the destination database is connected.</p> |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                               |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The character set of the source database is inconsistent with that of the destination database.</p> <p>Handling suggestion: Modify the character sets.</p> <p>Run commands to modify the self-created source database.</p> <ol style="list-style-type: none"> <li>1. Check whether source and destination database character sets are consistent.<br/> <code>show variables like "character_set_server"\G;</code><br/> <pre>mysql&gt; show variables like "character_set_server"\G; ***** 1. row ***** Variable_name: character_set_server Value: utf8 1 row in set (0.00 sec)</pre> </li> <li>2. Modify the character set of the source database server.<br/> <code>set character_set_server='utf8';</code><br/> <pre>mysql&gt; set character_set_server='utf8'; Query OK, 0 rows affected (0.00 sec)</pre> </li> </ol> |
|--|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 6.1.2.5 Checking Whether the Source Database server\_id Meets the Incremental Migration Requirements

#### MySQL

**Table 6-9** Checking whether the source database server\_id meets the incremental migration requirements

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database <b>server_id</b> meets the incremental migration requirements                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                           | Check whether the source database <b>server_id</b> meets the incremental migration requirements.                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database failed to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p>                                                                                                                                                                                                                                                                                                       |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                                                                                                                                                                                                                                                                                                                            |
|                                              | <p>Failure cause: The source database <b>server_id</b> does not meet the incremental migration requirements.</p> <p>Handling suggestion:</p> <p>Run the following command to modify the <b>server_id</b> value:<br/> <code>set global server_id=n</code></p> <p>The value <b>n</b> indicates the source database server_id. If the source database version is MySQL 5.6, the value <b>n</b> ranges from 2 to 4294967296. Otherwise, the value <b>n</b> ranges from 1 to 4294967296.</p> |

## 6.1.2.6 Checking Whether the Source and Destination Database Table Names Are Consistent in Case Sensitivity

### MySQL Migration

**Table 6-10** Checking whether the source and destination database table names are consistent in case sensitivity

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source and destination database table names are consistent in case sensitivity                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                           | Check whether the source and destination database names and table names are consistent in case sensitivity.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                                              | Failure cause: Insufficient user permissions<br>Handling suggestion: Check whether the database user permissions meet the migration requirements.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|                                              | Failure cause: The <b>lower_case_table_names</b> values in the source and destination databases must be the same.<br>Handling suggestion: <ul style="list-style-type: none"><li>• If you are migrating data out of the cloud, change the values of <b>lower_case_table_names</b> in the source and destination databases to the same.<br/>You are advised to change the parameter value in an empty database. For example, if the destination RDS DB instance is empty, run the following example command to change the <b>lower_case_table_names</b> value to the same as that in the source database: Sample command:<br/><pre>set global lower_case_table_names=n;</pre><br/>In the preceding command, <b>n</b> indicates the parameter value of the source database. After the modification, restart the database for the modification to take effect.</li><li>• If you are migrating data out of the cloud, perform the following operations:<br/>If the destination database is a self-built database, modify the <b>lower_case_table_names</b> parameter of the destination database. Add <b>lower_case_table_names=n</b> under the <b>[mysqld]</b> tag in the MySQL configuration file <b>my.cnf</b>. <b>n</b> indicates the value of parameter <b>same lower_case_table_names</b> of the source database. The database must be restarted to make the change take effect.<br/>If the destination database is a cloud database, check whether the <b>lower_case_table_names</b> parameter can be modified. If not, contact technical support.</li></ul> |

|  |                                                                                                                                                                                                                                                                                                                                                                         |
|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | Failure cause: The <b>lower_case_table_names</b> parameter value of the destination database is different from that of the source database, and the source database contains uppercase database and table names.<br>Handling suggestion: Rectify the fault by referring to <a href="#">How Can I Migrate Databases or Tables Whose Names Contain Uppercase Letters?</a> |
|  | Failure cause: The database is unavailable.<br>Handling suggestion: Contact technical support.                                                                                                                                                                                                                                                                          |

### 6.1.2.7 Checking Whether the Source Database Contains Object Names with Non-ASCII Characters

#### MySQL

**Table 6-11** Checking whether the source database contains object names with non-ASCII characters

|                                              |                                                                                                                                                                                                    |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database contains object names with non-ASCII characters                                                                                                                        |
| <b>Description</b>                           | If the source database contains object names with non-ASCII characters, the migration will fail.                                                                                                   |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The source database cannot contain object names with non-ASCII characters.<br>Handling suggestion: In the source database, change the object names containing non-ASCII characters. |

### 6.1.2.8 Checking Whether the TIME\_ZONE Values of the Source and Destination Databases Are the Same

#### MySQL

**Table 6-12** Checking whether the TIME\_ZONE values of the source and destination databases are the same

|                    |                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the <b>TIME_ZONE</b> values of the source and destination databases are the same                       |
| <b>Description</b> | The migration fails because the <b>TIME_ZONE</b> values of the source and destination databases are different. |



|                                              |                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The <b>TIME_ZONE</b> or <b>SYSTEM_TIME_ZONE</b> values of the source and destination databases must be the same.</p> <p>Handling suggestion: Change the <b>TIME_ZONE</b> value of the destination database to the same as that of the source database, or change the <b>TIME_ZONE</b> value of the source database to the same as that of the destination database.</p> |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 6.1.2.9 Checking Whether the COLLATION\_SERVER Values of the Source and Destination Databases Are the Same

#### MySQL

**Table 6-13** Checking whether the COLLATION\_SERVER values of the source and destination databases are the same

|                                              |                                                                                                                                                                                                                                          |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>COLLATION_SERVER</b> values of the source and destination databases are the same                                                                                                                                          |
| <b>Description</b>                           | The migration fails because the <b>COLLATION_SERVER</b> values of the source and destination databases are different.                                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The <b>COLLATION_SERVER</b> values of the source and destination databases must be the same.</p> <p>Handling suggestion: Change <b>COLLATION_SERVER</b> of the source and destination databases to the same value.</p> |

### 6.1.2.10 Checking Whether the SERVER\_UUID Values of the Source and Destination Databases Are the Same

#### MySQL Migration

**Table 6-14** Checking whether the SERVER\_UUID values of the source and destination databases are the same

|                    |                                                                                                             |
|--------------------|-------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the <b>SERVER_UUID</b> values of the source and destination databases are the same                  |
| <b>Description</b> | If the <b>SERVER_UUID</b> values of the source and destination databases are the same, the migration fails. |

|                                              |                                                                                                                                                                                                                  |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>SERVER_UUID</b> values of the source and destination databases must be different.<br>Handling suggestion: Check that the source and destination databases are not the same MySQL database. |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 6.1.2.11 Checking Whether the **SERVER\_ID** Values of the Source and Destination Databases Are Different

#### MySQL

**Table 6-15** Checking whether the **SERVER\_ID** values of the source and destination databases are different

|                                              |                                                                                                                                                                                                                    |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>SERVER_ID</b> values of the source and destination databases are different                                                                                                                          |
| <b>Description</b>                           | Check whether the <b>SERVER_ID</b> values of the source and destination databases are different. If they are the same, the migration fails.                                                                        |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>SERVER_ID</b> values of the source and destination databases must be different.<br>Handling suggestion: Change <b>SERVER_ID</b> of the source and destination databases to different values. |

### 6.1.2.12 Checking Whether the Source Database Contains Invalid **sql\_mode** Values

#### MySQL

**Table 6-16** Checking whether the source database contains invalid **sql\_mode** values

|                                              |                                                                                                                                                                                                    |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database contains invalid <b>sql_mode</b> values                                                                                                                                |
| <b>Description</b>                           | If the source database contains invalid <b>sql_mode</b> values, the migration will fail.                                                                                                           |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>sql_mode</b> value of the source database cannot be <b>no_engine_substitution</b> .<br>Handling suggestion: Change <b>sql_mode</b> of the source database to a proper value. |

### 6.1.2.13 Checking Whether the SQL\_MODE Values of the Source and Destination Databases Are the Same

#### MySQL

**Table 6-17** Checking whether the SQL\_MODE values of the source and destination databases are the same

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>SQL_MODE</b> values of the source and destination databases are the same                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                           | Check whether the <b>SQL_MODE</b> values of source and destination databases are the same. If they are inconsistent, the migration may fail.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Failure Cause and Handling Suggestion</b> | <ul style="list-style-type: none"><li>If you are migrating data to the cloud, perform the following operations:<br/>Failure cause: The <b>SQL_MODE</b> values of the source and destination databases must be the same.<br/>Handling suggestion: Change the <b>SQL_MODE</b> values of the destination database to the same as those of the source database, Ensure that both the source and destination databases do not have the forbidden SQL_MODE values. For details, see "Modifying Parameters in a Parameter Group". If MyISAM tables are to be migrated, the <b>SQL_MODE</b> values in the destination database cannot contain <b>NO_ENGINE_SUBSTITUTION</b>.</li><li>If you are migrating data out of the cloud, perform the following operations:<br/>Item to be confirmed: The <b>SQL_MODE</b> values of the source and destination databases must be the same.<br/>Handling suggestions: Change the <b>SQL_MODE</b> values of the destination database to the same as those of the source database. Ensure that both the source and destination databases do not have the forbidden <b>SQL_MODE</b> values.</li></ul> |

### 6.1.2.14 Checking Whether the sql\_mode Value in the Destination Database Is Not NO\_ENGINE\_SUBSTITUTION

#### MySQL Migration and Synchronization

**Table 6-18** Checking whether the sql\_mode value in the destination database is not NO\_ENGINE\_SUBSTITUTION

|                    |                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the sql_mode value in the destination database is not NO_ENGINE_SUBSTITUTION                                                  |
| <b>Description</b> | MyISAM tables are included in the migration objects. The sql_mode value in the destination database cannot be NO_ENGINE_SUBSTITUTION. |

|                                              |                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The <b>sql_mode</b> value in the destination database is <b>NO_ENGINE_SUBSTITUTION</b>.</p> <p>Handling suggestion: In the destination database, set <b>SQL_MODE</b> to a value other than <b>NO_ENGINE_SUBSTITUTION</b>. For details, see "Modifying Parameters" in the <i>Relational Database Service User Guide</i>.</p> |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 6.1.2.15 Checking Whether the innodb\_strict\_mode Values of the Source and Destination Databases Are the Same

#### MySQL Migration

**Table 6-19** Checking whether the innodb\_strict\_mode values of the source and destination databases are the same

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>innodb_strict_mode</b> values of the source and destination databases are the same                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b>                           | Check whether the <b>innodb_strict_mode</b> values of source and destination databases are the same. If they are inconsistent, the migration may fail.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | <ul style="list-style-type: none"> <li>If you are migrating data to the cloud, perform the following operations:<br/>Failure cause: The <b>innodb_strict_mode</b> values of the source and destination databases must be the same.<br/>Handling suggestion: Create a parameter group for the destination database and change <b>innodb_strict_mode</b> to the same value as that of the source database. For details, see <b>Creating a Parameter Group</b> in <i>Relational Database Service User Guide</i>.</li> <li>If you are migrating data out of the cloud, perform the following operations:<br/>Failure cause: The <b>innodb_strict_mode</b> values of the source and destination databases must be the same.<br/>Handling suggestion: Change <b>innodb_strict_mode</b> of the destination database to the same value as that of the source database.</li> </ul> |

## 6.1.2.16 Checking Whether the max\_wal\_senders Value of the Source Database Is Correctly Configured

### PostgreSQL Synchronization

**Table 6-20** Checking whether the max\_wal\_senders value of the source database is correctly configured

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>max_wal_senders</b> value of the source database is correctly configured                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b>                           | The <b>max_wal_senders</b> value of the source database must be greater than the number of used replication slots. Otherwise, the synchronization may fail.                                                                                                                                                                                                                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>max_wal_senders</b> value of the source database is less than or equal to the number of used replication slots.<br>Handling suggestion: Set <b>max_wal_senders</b> to a value greater than the number of used replication slots and restart the database to apply the changes. Run the following command to query the number of used replication slots in the current database:<br><pre>select count(1) from pg_replication_slots;</pre> |

## 6.1.2.17 Checking Whether the WAL\_LEVEL Value in the Source Database Is Correct

### PostgreSQL Synchronization

**Table 6-21** Checking whether the WAL\_LEVEL value in the source database is correct

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the WAL_LEVEL value in the source database is correct                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                           | Check whether <b>wal_level</b> of the source database is set to <b>logical</b> . If the value is not <b>logical</b> , the incremental logs of the source database cannot be logically decoded. As a result, incremental synchronization cannot be performed.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>wal_level</b> value in the source database is incorrect.<br>Handling suggestion: Change the <b>wal_level</b> value of the source database to <b>logical</b> . For details about how to modify the parameter for self-built databases, see: <ul style="list-style-type: none"><li>• Run <b>alter system set wal_level = logical</b> in the source database as a super user and restart the database to apply the changes.</li><li>• Alternatively, modify the <b>postgresql.conf</b> configuration file, set <b>wal_level</b> to <b>logical</b>, and restart the database to apply the changes.</li></ul> |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The source database version is not supported.</p> <p>Handling suggestion: Ensure that the source database version is supported by DRS. Supported source database versions include PostgreSQL 9.4, 9.5, 9.6, 10, 11, 12, and 13.</p>                                                                                                                                                                              |
|  | <p>Failure cause: The destination database version is not supported.</p> <p>Handling suggestion: Ensure that the destination database version is supported by DRS. The destination database supports the following major versions: RDS for PostgreSQL 9.5, 9.6, 10, 11, 12, and 13. If the source database is RDS for PostgreSQL Enhanced Edition, the destination database supports only RDS for PostgreSQL Enhanced Edition.</p> |

### 6.1.2.18 Checking Whether the MAX\_REPLICATION\_SLOTS Value in the Source Database Is Correct

#### PostgreSQL Synchronization

**Table 6-22** Checking whether the MAX\_REPLICATION\_SLOTS value in the source database is correct

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>MAX_REPLICATION_SLOTS</b> value in the source database is correct                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                           | The <b>max_replication_slots</b> value of the source database must be greater than the number of used replication slots. Otherwise, the synchronization may fail.                                                                                                                                                                                                                                                                                                                  |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The <b>max_replication_slots</b> value of the source database is less than or equal to the number of used replication slots.</p> <p>Handling suggestion: Set <b>max_replication_slots</b> to a value greater than the number of used replication slots and restart the database to apply the changes. Run the following command to query the number of used replication slots in the current database:</p> <pre>select count(1) from pg_replication_slots;</pre> |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the synchronization requirements.</p>                                                                                                                                                                                                                                                                                                                 |

## 6.1.2.19 Checking Whether the Source Database Is on Standby

### PostgreSQL Synchronization

**Table 6-23** Checking whether the source database is on standby

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database is on standby                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                           | <p>For a full+incremental synchronization task, the source database cannot be a standby database. Otherwise, incremental synchronization cannot be performed.</p> <p>For a full synchronization task, the source database can be a standby database, but <b>hot_standby_feedback</b> must be set to <b>on</b>. Otherwise, the synchronization may fail.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: In a real-time full+incremental synchronization task, the source database cannot be a standby database. Otherwise, incremental synchronization cannot be performed.</p> <p>Handling suggestion: Configure the source database as the primary database.</p> <p>Failure cause: For a full synchronization task, the source database <b>hot_standby_feedback</b> is set to <b>off</b>.</p> <p>Handling suggestion: The source database is configured as the primary database, or <b>hot_standby_feedback</b> of the source database is set to <b>on</b>.</p> <ul style="list-style-type: none"><li>• Change the source database to the primary database.</li><li>• Alternatively, change the <b>hot_standby_feedback</b> value of the source database to <b>on</b> before starting full synchronization. After the full synchronization is complete, change the value of this parameter to <b>off</b>.</li></ul> |

## 6.1.2.20 Checking Whether the log\_slave\_updates Value of the Source Database Is Correctly Configured

### MySQL Migration

**Table 6-24** Checking whether the log\_slave\_updates value of the source database is correctly configured

|                    |                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the <b>log_slave_updates</b> value of the source database is correctly configured             |
| <b>Description</b> | The migration will fail if the <b>log_slave_updates</b> parameter of the source database is disabled. |

|                                                     |                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b>        | Failure cause: The <b>slave_updates_check</b> parameter of the source database must be enabled.<br>Handling suggestion: In the MySQL configuration file <b>my.cnf</b> , add the "log_slave_updates=1" line under [mysqld] and restart the database for the modification to take effect.                                                               |
|                                                     | Failure cause: The source database is a standby database and the <b>log_slave_updates</b> value is <b>OFF</b> .<br>Handling suggestion: On the source database, set <b>log_slave_updates</b> to <b>ON</b> . Then, restart the database for the modification to take effect.                                                                           |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Item to be confirmed: The source database is a standby database and the <b>log_slave_updates</b> value is <b>OFF</b> .<br>Handling suggestion: On the source database, set <b>log_slave_updates</b> to <b>ON</b> . Then, restart the database for the modification to take effect. If no switchover or failover will occur, no operation is required. |

### 6.1.2.21 Checking Whether the BLOCK\_SIZE Value of the Source Database Is the Same as That of the Destination Database

#### PostgreSQL Synchronization

**Table 6-25** Checking whether the BLOCK\_SIZE value of the source database is the same as that of the destination database

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>BLOCK_SIZE</b> value of the source database is the same as that of the destination database                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                           | The <b>BLOCK_SIZE</b> value of the destination database must be greater than or equal to that of the source database. Otherwise, the synchronization may fail.                                                                                                                                                                                                                                                                                              |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>BLOCK_SIZE</b> value of the destination database is less than that of the source database.<br>Handling suggestion: <ul style="list-style-type: none"> <li>Use the destination database whose <b>BLOCK_SIZE</b> value is greater than or equal to that of the source database.</li> <li>Use the source database whose <b>BLOCK_SIZE</b> value is less than or equal to the value of destination database <b>BLOCK_SIZE</b>.</li> </ul> |



### 6.1.2.22 Checking Whether the binlog\_row\_image Value is FULL

#### MySQL

**Table 6-26** Checking whether the binlog\_row\_image value is FULL

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>binlog_row_image</b> value is <b>FULL</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                           | If the <b>binlog_row_image</b> value of the source database is not <b>FULL</b> , the migration will fail.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The <b>binlog_row_image</b> value of the source database is not <b>FULL</b>.</p> <p>Handling suggestion:</p> <ul style="list-style-type: none"><li>• If the source database is an RDS DB instance on the cloud, change <b>binlog_row_image</b> to <b>FULL</b> on the RDS console, and then restart the source database.</li><li>• If the source database is an on-premises database, perform the following steps:<ol style="list-style-type: none"><li>1. Log in to the server where the MySQL source database is located.</li><li>2. Manually change the value of <b>binlog_row_image</b> in the <b>my.cnf</b> configuration file to <b>FULL</b> and save the file.<pre>binlog_row_image=full</pre></li><li>3. To ensure a successful task, restart the source database during off-peak hours.</li></ol></li></ul> |

### 6.1.2.23 Checking Whether the Transaction Isolation Levels are Consistent

#### MySQL

**Table 6-27** Checking whether the transaction isolation levels are consistent

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the transaction isolation levels are consistent                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                           | Check whether the transaction isolation levels of the source and destination databases are the same.                                                                                                                                                                                                                                                                                                 |
| <b>Failure Cause and Handling Suggestion</b> | <p>If you are migrating data to the cloud, perform the following operations:</p> <p><b>Failure cause:</b> The transaction isolation levels of the source and destination databases are different.</p> <p><b>Handling suggestion:</b> Change the isolation level (<b>tx_isolation</b> or <b>transaction_isolation</b>) of the destination database to be the same as that of the source database.</p> |

## 6.1.2.24 Checking Whether the lc\_monetary Values of the Source and Destination Databases Are the Same

### PostgreSQL Synchronization

**Table 6-28** Checking whether the lc\_monetary values of the source and destination databases are the same

|                                              |                                                                                                                                                                                                                                                           |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the <b>lc_monetary</b> values of the source and destination databases are the same                                                                                                                                                                |
| <b>Description</b>                           | Check whether the <b>lc_monetary</b> values of the source and destination databases are the same. If they are inconsistent, the synchronization fails.                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database failed to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                                                    |
|                                              | Failure cause: This item cannot be checked because the destination database failed to be connected.<br>Handling suggestion: Check whether the destination database is connected.                                                                          |
|                                              | Failure cause: The <b>lc_monetary</b> values of the source and destination databases must be the same.<br>Handling suggestion: Check whether the <b>lc_monetary</b> values of the source and destination databases meet the synchronization requirements. |
|                                              | Failure cause: Insufficient user permissions<br>Handling suggestion: Check whether the database user permissions meet the synchronization requirements.                                                                                                   |

## 6.1.2.25 Checking Whether the Source Database Contains Trigger Names with Non-ASCII Characters

### MySQL

**Table 6-29** Checking whether the source database contains trigger names with non-ASCII characters

|                    |                                                                                |
|--------------------|--------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the source database contains trigger names with non-ASCII characters   |
| <b>Description</b> | If the source database contains non-ASCII characters, the migration will fail. |

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed: The source database cannot contain view names with non-ASCII characters.</p> <p>Handling suggestion: To solve this problem, perform the following steps:</p> <p>Method 1:</p> <p>Click <b>Previous</b> to return to the <b>Select Migration Type</b> page. Select a customized object and do not select the trigger name that contains non-ASCII characters.</p> <p>Method 2: Change the trigger name.</p> |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 6.1.2.26 Checking Whether the Source Database Collections Contain More Than 10 Indexes

#### MongoDB Migration

**Table 6-30** Checking whether the source database collections contain more than 10 indexes

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database collections contain more than 10 indexes                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                           | If the number of indexes in the source database exceeds 10, the migration duration is affected.                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Failure Cause and Handling Suggestion</b> | <p>Alarm cause: The source database has collections containing more than 10 indexes, which are migrated slowly.</p> <p>Handling suggestion: The number of indexes affects the migration duration. Check whether all indexes need to be migrated. If the index does not need to be migrated, delete the index before starting the migration.</p> <p>Run the following command to delete the index:<br/>run the <code>db.Collection name.dropIndex(Index name)</code></p> |

### 6.1.2.27 Checking Whether log\_bin\_trust\_function\_creators Is Set to On in Both the Source and Destination Databases

#### MySQL

**Table 6-31** Checking whether log\_bin\_trust\_function\_creators is set to on in both the source and destination databases

|                   |                                                                                                   |
|-------------------|---------------------------------------------------------------------------------------------------|
| <b>Check Item</b> | Whether log_bin_trust_function_creators is set to on in both the source and destination databases |
|-------------------|---------------------------------------------------------------------------------------------------|

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>                                  | During the out-of-cloud migration from MySQL to MySQL, the <b>log_bin_trust_function_creators</b> value of the source database must be the same as that of the destination database. If the source database supports user-defined functions (UDFs) but the destination database does not, change the <b>log_bin_trust_function_creators=off</b> parameter of the destination database to <b>log_bin_trust_function_creators=on</b> . If the parameters of the source and destination are different, the migration may fail. |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Item to be confirmed: The destination database does not support custom functions.<br>Handling suggestions: In the <b>my.cnf</b> file of the destination database, check whether <b>log_bin_trust_function_creators=on</b> exists. If it does not exist, add <b>log_bin_trust_function_creators=on</b> and restart the database for the modification to take effect.                                                                                                                                                         |

### 6.1.2.28 Checking Whether log\_bin\_trust\_function\_creators Is Set to On in the Destination Database

#### MySQL

**Table 6-32** Checking whether log\_bin\_trust\_function\_creators is set to on in the destination database

|                                              |                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether <b>log_bin_trust_function_creators</b> is set to <b>on</b> in the destination database                                                                                                                                                                                                                                                               |
| <b>Description</b>                           | During the migration from RDS for MySQL to MySQL out of the cloud, the destination database does not support custom functions.                                                                                                                                                                                                                               |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The destination database does not support custom functions.<br>Handling suggestions: In the <b>my.cnf</b> file of the destination database, check whether <b>log_bin_trust_function_creators=on</b> exists. If it does not exist, add <b>log_bin_trust_function_creators=on</b> and restart the database for the modification to take effect. |

### 6.1.2.29 Checking Whether the max\_allowed\_packet Value of the Destination Database Is too Small

#### MySQL Migration

**Table 6-33** Checking whether the max\_allowed\_packet value of the destination database is too small

|                                              |                                                                                                                                                                                                                                                  |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the max_allowed_packet value of the destination database is too small                                                                                                                                                                    |
| <b>Description</b>                           | A large amount of data cannot be written to the destination database during the migration because the max_allowed_packet value is smaller than 100 MB. As a result, the full migration failed.                                                   |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The <b>max_allowed_packet</b> value of the destination database is too small, which may cause data fails to be written during the migration.<br>Handling suggestions: Set the <b>max_allowed_packet</b> value greater than 100 MB |

### 6.1.2.30 Checking Whether the Databases and Tables Exist

#### All Scenarios

**Table 6-34** Checking whether the databases and tables exist

|                                              |                                                                                                                                                                           |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the databases and tables exist                                                                                                                                    |
| <b>Description</b>                           | There are databases and tables in the uploaded file that do not exist in the source database. The synchronization fails.                                                  |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: Objects imported from files do not exist in the source database.<br>Handling suggestion: Remove these objects that do not exist and import the file again. |

### 6.1.2.31 Checking Whether session\_replication\_role of the Destination Database Is correctly Set

#### PostgreSQL Synchronization

**Table 6-35** Checking whether the session\_replication\_role value of the destination database is correctly set

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the <b>session_replication_role</b> value of the destination database is correctly set.                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                  | The <b>session_replication_role</b> parameter of the destination database is not set to <b>replica</b> . Data synchronization may fail when the synchronized table has associated foreign key constraints or triggers.                                                                                                                                                                                                                                                             |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Item to be confirmed: The <b>session_replication_role</b> parameter of the destination database is not set to <b>replica</b> .<br>Handling suggestion: Before starting the synchronization task, set <b>session_replication_role</b> of the destination database to <b>replica</b> . After the synchronization is complete, change the value of this parameter to <b>origin</b> . If the destination database is an RDS instance, you can modify the parameter on the RDS console. |

### 6.1.2.32 Checking Whether the MongoDB Instance Type Matches the Migration Mode

#### MongoDB Migration

**Table 6-36** Checking whether the MongoDB instance type matches the migration mode

|                                              |                                                                                                                                                                                                                                                                                      |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the MongoDB instance type matches the migration mode                                                                                                                                                                                                                         |
| <b>Description</b>                           | Check whether the MongoDB instance type matches the migration mode. If not, the migration fails.                                                                                                                                                                                     |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: When a DRS task is created, the source DB instance type is set to <b>Cluster</b> , but the source database is not a cluster.<br>Handling suggestion: If the source DB instance type is set to <b>Cluster</b> , ensure that the source database is a cluster database. |

### 6.1.2.33 Checking the Physical Standby Database

#### Oracle Synchronization

**Table 6-37** Physical standby database check

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Physical standby database check                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>                                  | If the source Oracle database is in the incremental migration phase, the instance type of the source must be the same as that of the destination. Otherwise, the migration fails.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed:</p> <ol style="list-style-type: none"> <li>1. The physical standby database does not generate logs. It replicates them from the primary database. Check whether supplemental logging of the primary database meets the incremental synchronization requirements.<br/>Handling suggestion: For details, see <a href="#">How Do I Check Supplemental Logging of the Source Oracle Database?</a></li> <li>2. The physical standby database does not generate logs, resulting in synchronization task delay. You can shorten the interval for archiving logs from the primary database to the physical standby database. However, extremely low values can result in a large number of logs, so you are advised to synchronize data from the logical standby database. Run the following statement on the primary database to specify the log archive interval:<br/><code>alter system set archive_lag_target=seconds;</code></li> </ol> |

### 6.1.2.34 Checking Whether the Values of group\_concat\_max\_len Are Consistent

#### MySQL Migration

**Table 6-38** group\_concat\_max\_len consistency check

|                    |                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | The <b>group_concat_max_len</b> value in the destination database is inconsistent with that in the source database.                                                               |
| <b>Description</b> | If the values of <b>group_concat_max_len</b> in the source and destination databases are different, the queried fields may be truncated. Change the parameter values to the same. |

|                                                     |                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed: If the values of group_concat_max_len in the source and destination databases are different, the queried fields may be truncated. Change the parameter values to the same.</p> <p>Handling suggestion: Change the parameter values to the same.</p> |
|-----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 6.1.2.35 Checking Whether the Character Sets Are Compatible

#### Oracle Synchronization

**Table 6-39** Character set compatibility check

|                                                     |                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Character set compatibility check                                                                                                                                                                                                                            |
| <b>Description</b>                                  | The character set of the destination database is incompatible with that of the source database.                                                                                                                                                              |
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed: The character set of the destination database is incompatible with that of the source database.</p> <p>Handling suggestion: Change the character set of the destination database to be the same as that of the source database.</p> |

### 6.1.2.36 Checking Replication Attribute of Primary Key Columns

#### PostgreSQL as the Source

**Table 6-40** Replication attribute check of primary key columns

|                    |                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Replication attribute check of primary key columns                                                                                                                          |
| <b>Description</b> | During a full+incremental synchronization or an incremental synchronization task, the replication attribute of primary key columns in the source database table is checked. |



|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed: All primary key columns of the tables to be synchronized are columns whose storage attribute is plain, and the replication attribute of the tables is neither full nor default. Incremental synchronization may fail.</p> <p>Handling suggestion: Run the following SQL statement to change the replication attribute of the tables to default:<br/>alter table schema.table replica identity default;</p>                                                                                                                                       |
|                                                     | <p>Item to be confirmed: The primary key columns of the tables to be synchronized contain columns whose storage attribute is not plain, and the replication attribute of the tables is neither full nor default. There is a high probability that incremental synchronization will fail.</p> <p>Handling suggestion: Run the following SQL statement to change the replication attribute of the tables to full: (If the replication attribute is changed to default, incremental synchronization may still fail.)<br/>alter table schema.table replica identity full;</p> |
|                                                     | <p>Item to be confirmed: The primary key columns of the tables to be synchronized contain columns whose storage attribute is not plain, and the replication attribute of the tables is not full. Incremental synchronization may fail.</p> <p>Handling suggestion: Run the following SQL statement to change the replication attribute of the tables to full:<br/>alter table schema.table replica identity full;</p>                                                                                                                                                     |

### 6.1.2.37 Checking Whether the Source and Destination Database Character Sets Are Consistent

#### MySQL->MySQL Migration

**Table 6-41** Checking whether the source and destination database character sets are consistent

|                                                     |                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the source and destination database character sets are consistent                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | Checking whether the source and destination database character sets are consistent                                                                                                                                                                                                                                                  |
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed: The source database supports character sets of a later version.</p> <p>Handling suggestion: The source database supports character sets of a later version. Check whether the source database uses a character set of a later version and whether the destination database supports the character set.</p> |

## 6.1.3 Source DB Instance Statuses

### 6.1.3.1 Checking Whether the Source and Destination Databases Are of the Same Type

#### MongoDB Migration

**Table 6-42** Checking whether the source and destination databases are of the same type

|                                              |                                                                                                                                                                               |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source and destination databases are of the same type                                                                                                             |
| <b>Description</b>                           | Check whether the source and destination databases are of the same type. If they are of different types, the migration fails.                                                 |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The destination database is a cluster but the source database is a replica set.<br>Handling suggestion: Change the type of the source or destination database. |
|                                              | Failure cause: The destination database is a replica set but the source database is a cluster.<br>Handling suggestion: Change the type of the source or destination database. |

### 6.1.3.2 Checking Whether the ChangeStream API of the source DB instance is available

#### MongoDB-to-DDS Migration

**Table 6-43** Checking whether the ChangeStream API of the source DB instance is available

|                    |                                                                         |
|--------------------|-------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the ChangeStream API of the source DB instance is available     |
| <b>Description</b> | Check whether the ChangeStream API of the source database is available. |

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The source database cannot use the ChangeStream API.<br>Handling suggestion:<br><ol style="list-style-type: none"><li>1. Check whether the source database version is MongoDB 4.0 or later.</li><li>2. Check whether the WiredTiger storage engine is enabled for the source database. If not, you are advised to create a DRS task and select the oplog mode. Run the following command (on a shard node):<br/><pre>db.serverStatus().storageEngine.name;</pre></li></ol> |
|                                              | Failure cause: The source database cannot use the ChangeStream API.<br>Handling suggestion: Check whether the source database version is DDS 4.0 or later. If not, upgrade the source database to DDS 4.0 or later.                                                                                                                                                                                                                                                                       |

## 6.1.4 Destination DB Instance Statuses

### 6.1.4.1 Checking Whether the Destination Database Is Involved in Another Migration Task

#### MySQL

**Table 6-44** Checking whether the destination database is involved in another migration task

|                                              |                                                                                                                                                                                                                |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database is involved in another migration task                                                                                                                                         |
| <b>Description</b>                           | Check whether the destination database is being used in another migration task. If more than one migration task uses the same destination database, the migration may fail.                                    |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The destination RDS DB instance is being used in another migration task.<br>Handling suggestion: Wait for the migration task to complete. You can also stop or delete an unused migration task. |

## 6.1.4.2 Checking Whether the Destination Database Has a Read Replica

### MySQL

**Table 6-45** Checking whether the destination database has a read replica

|                                              |                                                                                                                                                                                                                                        |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database has a read replica                                                                                                                                                                                    |
| <b>Description</b>                           | Check whether the destination database has read replicas. If the destination database has read replicas, the incremental migration may fail.                                                                                           |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: In an incremental migration, the destination database cannot have read replicas.<br>Handling suggestion: Delete the read replicas from the destination database. After the migration is complete, create read replicas. |

## 6.1.4.3 Checking Whether the Extensions Are Supported

### PostgreSQL Synchronization

**Table 6-46** Checking whether the extensions are supported

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the extensions are supported                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>                           | Check whether the source database has plug-ins that are not installed on the destination database.                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: Extensions installed in the source database are not supported in the destination database.<br>Handling suggestion: <ul style="list-style-type: none"><li>• If the source database services do not depend on those extensions, run the following statement to delete the extensions. Replace <i>plugin_name</i> with the name of the extension to be deleted.<br/>drop extension plugin_name;</li><li>• Alternatively, use a destination database that supports these extensions.</li></ul> |
|                                              | Failure cause: The source database has extensions that contain tables as members.<br>Handling suggestion: Check whether the source database extensions contain metadata generated after the extensions are created. If yes, use the dedicated syntax of the extension to rebuild the metadata after the migration is complete.                                                                                                                                                                            |

|  |                                                                                                                                                                                                                                                                                                                                                                               |
|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The destination database user does not have the permission to create extensions.</p> <p>Handling suggestion: Grant the permission to the user in the destination database as user root. Run the following SQL statements (replace <i>username</i> with the destination database username):<br/>alter user username inherit;<br/>grant root to username;</p> |
|  | <p>Failure cause: The extension version supported by the destination database is earlier than that installed in the source database.</p> <p>Handling suggestion: Use the destination database that supports extensions of a later version (not earlier than the source database extension version) and create a synchronization task again.</p>                               |

### 6.1.4.4 Checking Whether the Destination DB Instance Is Available

**Table 6-47** Checking whether the destination DB instance is available

|                                              |                                                                                                                                                                          |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination DB instance is available                                                                                                                         |
| <b>Description</b>                           | Check whether the primary instance and read replicas are available in the destination database. If not, the migration fails.                                             |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The destination DB instance is not available.</p> <p>Handling suggestion: Repair the destination DB instance.</p>                                      |
|                                              | <p>Failure cause: Read replicas in the destination database are abnormal.</p> <p>Handling suggestion: Repair the abnormal read replicas in the destination database.</p> |
|                                              | <p>Failure cause: The RDS service is abnormal. Try again later.</p> <p>Handling suggestion: Try again later.</p>                                                         |

## 6.1.5 Database User Permissions

### 6.1.5.1 Whether the Source Database User Has Sufficient Permissions

#### MySQL Migration

**Table 6-48** Whether the source database user has sufficient permissions

|                   |                                                             |
|-------------------|-------------------------------------------------------------|
| <b>Check Item</b> | Whether the source database user has sufficient permissions |
|-------------------|-------------------------------------------------------------|

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>                           | <p>The source database user must have required permissions for full and incremental migrations.</p> <ul style="list-style-type: none"><li>• In a full migration, the source database user must have the SELECT, SHOW VIEW, and EVENT permissions.</li><li>• In an incremental migration, the source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, and REPLICATION CLIENT.</li></ul> <p>If the permissions are insufficient, the migration will fail.</p> |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: In a full migration, the source database user must have the SELECT, SHOW VIEW, and EVENT permissions.</p> <p>Handling suggestions: Grant the source database user the corresponding permissions.</p>                                                                                                                                                                                                                                                                                                  |
|                                              | <p>Failure cause: In an incremental migration, the source database user must have the following permissions: SELECT, SHOW VIEW, EVENT, LOCK TABLES, REPLICATION SLAVE, REPLICATION CLIENT, In the DR scenario, the following permissions are required: CREATE, ALTER, DROP, DELETE, INSERT, UPDATE, and INDEX.</p> <p>Handling suggestions: Grant the source database user the corresponding permissions.</p>                                                                                                           |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p> <p><b>NOTE</b><br/>For details about the required MySQL permissions and authorized operations, see <a href="#">Which MySQL Permissions Are Required for DRS?</a></p>                                                                                                                                                                                       |

## PostgreSQL Synchronization

**Table 6-49** Checking whether the source database user has sufficient permissions

|                   |                                                             |
|-------------------|-------------------------------------------------------------|
| <b>Check Item</b> | Whether the source database user has sufficient permissions |
|-------------------|-------------------------------------------------------------|

|                                                                                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Descri<br/>ption</b></p>                                                     | <p>Specific permissions are required for different synchronization task types.</p> <ul style="list-style-type: none"> <li>• <b>Full synchronization:</b> CONNECT permission for databases, USAGE permission for schemas, SELECT permission for tables, SELECT permission for sequences, and SELECT permission for system table <b>catalog pg_catalog.pg_authid</b> (used for synchronizing user passwords)</li> <li>• <b>Full+incremental synchronization:</b> CONNECT permission for databases, USAGE permission for schemas, SELECT permission for tables, SELECT permission for sequences, SELECT permission for system table <b>pg_catalog.pg_authid</b> (used for synchronizing user passwords), UPDATE, DELETE, and TRUNCATE permissions for tables that do not have primary keys, and the permission to create a replication connection</li> </ul> <p>If the permissions are insufficient, the migration will fail.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <p><b>Failur<br/>e<br/>Cause<br/>and<br/>Handl<br/>ing<br/>Sugge<br/>stion</b></p> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.<br/>Handling suggestion: Check whether the source database is connected.</p> <p>Failure cause: Insufficient user permissions<br/>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p> <p>Failure cause: In a full migration, the source database user must have the SELECT, REFERENCES, TRIGGER, EXECUTE, and USAGE permissions.<br/>Handling suggestion: Change or authorize the migration account.</p> <p>Failure cause: The replication permission is not configured in <b>pg_hba.conf</b> for the replication instance and database user.<br/>Handling suggestion:<br/>Grant the replication permission to the user.<br/>Add the following to <b>pg_hba.conf</b>, and restart the database for the modification to take effect:<br/><b>host replication XXX(dbuser) 0.0.0.0/0 method</b><br/>After the migration is complete, delete this record and restart the database again.</p> <p>Failure cause: The <b>max_wal_senders</b> value in the source database is too small.<br/>Handling suggestion: In the <b>postgresql.conf</b> file, change the <b>max_wal_senders</b> value to a larger one, such as increasing it by 5 or 10.</p> <p>Failure cause: The database is unavailable.<br/>Handling suggestion: Contact technical support.</p> |

|                                                     |                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Item to Be Confirmed and Handling Suggestion</b> | <p>Item to be confirmed: The source database contains objects that can only be created by a superuser. The destination user is not a superuser, so the objects will be ignored.</p> <p>Handling suggestion: Use a superuser of the destination database or confirm that these objects can be ignored.</p> |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## MongoDB Migration

**Table 6-50** Checking whether the source database user has sufficient permissions

|                                              |                                                                                                                                                                                                                                                                                                                                                               |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database user has sufficient permissions                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                           | Check whether the source database user permissions meet the migration requirements. If the permissions are insufficient, the migration will fail.                                                                                                                                                                                                             |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p>                                                                                                                                                                              |
|                                              | <p>Failure cause: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database.</p> <p>Handling suggestion: Grant permissions to the source database user.<br/> <code>db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"}, {role:"read",db:"local"}])</code></p>   |
|                                              | <p>Failure cause: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database.</p> <p>Handling suggestion: Grant permissions to the source database user.<br/> <code>db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"}, {role:"read",db:"config"}])</code></p> |



|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The source mongos node user must have the readAnyDatabase permission for the admin database, and the read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database, and the read permission for the local database.</p> <p>Handling suggestion: Grant permissions to the source database user.</p> <ul style="list-style-type: none"> <li>• Statement for mongos:<br/>db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"},{role:"read",db:"config"}])</li> <li>• Statement for shard:<br/>db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"},{role:"read",db:"local"}])</li> </ul> |
|  | <p>Failure cause: The source database user must have the readAnyDatabase permission for the admin database.</p> <p>Handling suggestion: Grant permissions to the source database user.<br/>db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"}])</p>                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|  | <p>Failure cause: The source database user must have the read permission for the local database.</p> <p>Handling suggestion: Grant permissions to the source database user.<br/>db.grantRolesToUser("Username",[{role:"read",db:"local"}])</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|  | <p>Failure cause: The source database user must have the read permission for the config database.</p> <p>Handling suggestion: Grant permissions to the source database user.<br/>db.grantRolesToUser("Username",[{role:"read",db:"config"}])</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

### 6.1.5.2 Checking Whether the Destination Database User Has Sufficient Permissions

#### MySQL Migration

**Table 6-51** Checking whether the destination database user has sufficient permissions

|                    |                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the destination database user has sufficient permissions                                                                                       |
| <b>Description</b> | Check whether the destination database user permissions meet the migration requirements. If the permissions are insufficient, the migration will fail. |

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The destination database user must have the following permissions: SELECT, CREATE, DROP, DELETE, INSERT, UPDATE, INDEX, EVENT, CREATE VIEW, CREATE ROUTINE, TRIGGER, and WITH GRANT OPTION. If the destination database version is in the range from 8.0.14 to 8.0.18, the SESSION_VARIABLES_ADMIN permission is required.</p> <p>Handling suggestion: Check whether the destination database user permissions meet the migration requirements.</p> |
|                                              | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                                                                                                                                                                                                                                                                                                          |

## PostgreSQL Synchronization

**Table 6-52** Checking whether the destination database user has sufficient permissions

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database user has sufficient permissions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                           | <p>Different permissions are granted based on the scope of objects to be synchronized.</p> <ul style="list-style-type: none"> <li>● <b>Database-level synchronization objects:</b><br/>The CREATEDB permission and the <b>root</b> user or a member of root are required (only in special scenarios). For details, see the following description.</li> <li>● <b>Table-level synchronization object:</b><br/>To synchronize databases: The CREATEDB permission is required.<br/>To synchronize schemas, the CONNECT and CREATE permissions for the database that contains the schema are required.<br/>To synchronize objects in a schema: The CONNECT permission for the database that contains the schema, the USAGE permission for the schema that contains the object, and CREATE permission for the schema that contains the object are required.</li> <li>● <b>Synchronization user:</b> The CREATEROLE permission is required.</li> <li>● <b>Permission to synchronize objects:</b> The default privilege is the default value and cannot be modified. Otherwise, the object permissions of the destination database may be different from those of the source database.</li> </ul> |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The destination database user must have the CREATEDB permission.</p> <p>Handling suggestion: Grant the destination database user the CREATEDB permission.</p> <pre>alter role username with createdb;</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

|  |                                                                                                                                                                                                                                   |
|--|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The user does not have the USAGE permission for schemas.</p> <p>Handling suggestion: Grant the destination database user the CREATEDB permission.<br/>grant usage on schema schemaname to username;</p>         |
|  | <p>Failure cause: The user does not have the permission to create tables.</p> <p>Handling suggestion: Grant the destination database user the permission to create tables.<br/>grant create on schema schemaname to username;</p> |
|  | <p>Failure cause: The user does not have the permission to create schemas.</p> <p>Handling suggestion: Grant the destination database user the permission to create schemas.<br/>grant create on database to username;</p>        |

## MongoDB Migration

**Table 6-53** Checking whether the destination database user has sufficient permissions

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database user has sufficient permissions                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>                           | Check whether the destination database user permissions meet the migration requirements. If the permissions are insufficient, the migration will fail.                                                                                                                                                                                                                                                                                                                                           |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the destination database fails to be connected.</p> <p>Handling suggestion: Check whether the destination database is connected.</p>                                                                                                                                                                                                                                                                                                       |
|                                              | <p>Failure cause: The destination database user must have the dbAdminAnyDatabase permission for the admin database, the read permission for the config database, and the readWrite permission for the destination database.</p> <p>Handling suggestion: Run the following statement to grant permissions to the destination database user.<br/>db.grantRolesToUser("Username",[{role:"dbAdminAnyDatabase",db:"admin"}, {role:"read",db:"config"}, {role:"readWriteAnyDatabase",db:"admin"}])</p> |
|                                              | <p>Failure cause: The destination database user must have the readWrite permission for the databases to be migrated.</p> <p>Handling suggestion: Run the following statement to grant permissions to the destination database user.<br/>db.grantRolesToUser("Username",[{role:"readWriteAnyDatabase",db:"admin"}])</p>                                                                                                                                                                           |

|  |                                                                                                                                                                                                                                                                                                                            |
|--|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The destination database user must have the read permission for the config database.</p> <p>Handling suggestion: Run the following statement to grant permissions to the destination database user.</p> <pre>db.grantRolesToUser("Username",[{role:"read",db:"config"}])</pre>                           |
|  | <p>Failure cause: The destination database user must have the dbAdminAnyDatabase permission for the admin database.</p> <p>Handling suggestion: Run the following statement to grant permissions to the destination database user.</p> <pre>db.grantRolesToUser("Username",[{role:"dbAdminAnyDatabase",db:"admin"}])</pre> |
|  | <p>Failure cause: The destination database user must have the clusterManager permission for the admin database.</p> <p>Handling suggestion: Run the following statement to grant permissions to the destination database user.</p> <pre>db.grantRolesToUser("Username",[{role:"clusterManager",db:"admin"}])</pre>         |

## 6.1.6 Database Versions

### 6.1.6.1 Checking Whether the Source Database Version Is Supported

#### MySQL Migration

**Table 6-54** Checking whether the source database version is supported

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database version is supported                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>                           | Check whether the source database version is MySQL 5.5.x, MySQL 5.6.x, and MySQL 5.7.x.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The source database version is not supported.</p> <p>Handling suggestion: Export data and then import it following the instructions provided in the "Migrating MySQL Data Using mysqldump" section in the <i>Relational Database Service User Guide</i>.</p> <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p> <p>Failure cause: This item cannot be checked because the source database fails to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p> |

## Migration from MongoDB to DDS

**Table 6-55** Checking whether the source database version is supported

|                                              |                                                                                                                                                                                                                                  |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database version is supported                                                                                                                                                                                 |
| <b>Description</b>                           | Check whether the source database version is MongoDB 3.2.x, 3.4.x, or 4.0.x.                                                                                                                                                     |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The source database version is not supported.<br>Handling suggestion: Check whether the source database version is MongoDB 3.2x, 3.4.x, and 4.0.x.                                                                |
|                                              | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                            |
|                                              | Failure cause: The incremental data is obtained in changeStream mode but the source database version is earlier than 4.0.<br>Handling suggestion: In changeStream mode, ensure that the source database version is 4.0 or later. |

### 6.1.6.2 Checking Whether the Destination Database Version Is Supported

#### MySQL Migration

**Table 6-56** Checking whether the destination database version is supported

|                                              |                                                                                                                                                                                                                                                                           |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database version is supported                                                                                                                                                                                                                     |
| <b>Description</b>                           | Check whether the destination database version is MySQL 5.6.x, or MySQL 5.7.x.                                                                                                                                                                                            |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The destination database version is not supported.<br>Handling suggestion: Export data and then import it following the instructions provided in the "Migrating MySQL Data Using mysqldump" section in the <i>Relational Database Service User Guide</i> . |
|                                              | Failure cause: Insufficient user permissions<br>Handling suggestion: Check whether the database user permissions meet the migration requirements.                                                                                                                         |
|                                              | Failure cause: This item cannot be checked because the destination database fails to be connected.<br>Handling suggestion: Check whether the destination database is connected.                                                                                           |

## MongoDB to DDS Migration

**Table 6-57** Checking whether the destination database version is supported

|                                              |                                                                                                                                                                                 |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database version is supported                                                                                                                           |
| <b>Description</b>                           | Check whether the destination database version is MongoDB 3.4.x, 4.0.x, or 4.2.x.                                                                                               |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The destination database version is not supported.<br>Handling suggestion: Check whether the destination database version is MongoDB 3.4.x, 4.0.x, or 4.2.x.     |
|                                              | Failure cause: This item cannot be checked because the destination database fails to be connected.<br>Handling suggestion: Check whether the destination database is connected. |

### 6.1.6.3 Checking Whether the Migration Is from an Earlier Database Version to the Same or a Later Version

#### MySQL Migration

**Table 6-58** Checking whether the migration is from an earlier database version to the same or a later version

|                                              |                                                                                                                                                                                                                       |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the migration is from an earlier database version to the same or a later version                                                                                                                              |
| <b>Description</b>                           | Check whether the source database version is earlier than or the same as the destination database version.                                                                                                            |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                 |
|                                              | Failure cause: This item cannot be checked because the destination database fails to be connected.<br>Handling suggestion: Check whether the destination database is connected.                                       |
|                                              | Failure cause: The source database version is not supported.<br>Handling suggestion: Check whether the source database version is supported. Currently, only MySQL 5.5.x, MySQL 5.6.x, and MySQL 5.7.x are supported. |

|  |                                                                                                                                                                                                                                                    |
|--|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The destination database version is not supported.<br/>Handling suggestion: Check whether the destination database version is supported. Currently, only MySQL 5.6.x, and MySQL 5.7.x are supported.</p>                         |
|  | <p>Failure cause: Insufficient user permissions<br/>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                                                                                          |
|  | <p>Failure cause: The destination database version must be the same as or later than the source database version.<br/>Handling suggestion: Ensure the source database version is earlier than or the same as the destination database version.</p> |

## PostgreSQL Synchronization

**Table 6-59** Checking whether the migration is from an earlier database version to the same or a later version

|                                              |                                                                                                                                                                                                                                               |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the migration is from an earlier database version to the same or a later version                                                                                                                                                      |
| <b>Description</b>                           | Check whether the source database version is earlier than or the same as the destination database version.                                                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database fails to be connected.<br/>Handling suggestion: Check whether the source database is connected.</p>                                                                 |
|                                              | <p>Failure cause: This item cannot be checked because the destination database fails to be connected.<br/>Handling suggestion: Check whether the destination database is connected.</p>                                                       |
|                                              | <p>Failure cause: Insufficient user permissions<br/>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                                                                                     |
|                                              | <p>Failure cause: The source database version is not supported.<br/>Handling suggestion: Check whether the source database version is supported. Currently, only PostgreSQL 9.4, 9.5, 9.6, 10.0, 11.0, 12.0, and 13.0 are supported.</p>      |
|                                              | <p>Failure cause: The destination database version is not supported.<br/>Handling suggestion: Check whether the destination database version is supported. Currently, only PostgreSQL 9.5, 9.6, 10.0, 11.0, 12.0, and 13.0 are supported.</p> |

|  |                                                                                                                                                                                                                                                                                                                                                |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The major versions of the source and destination databases must be the same and the minor version of the source database must be less than or equal to that of the destination database.</p> <p>Handling suggestion: Ensure the source database version is earlier than or the same as the destination database version.</p> |
|  | <p>Failure cause: The destination database version and source database version do not meet the requirements of the selected migration mode.</p> <p>Handling suggestion: Check whether the versions of the destination database and source database meet the migration mode requirements.</p>                                                   |

## MongoDB Migration

**Table 6-60** Checking Whether the Source Database Version Is Supported

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the migration is from an earlier database version to the same or a later version                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                           | Check whether the source database version is earlier than or the same as the destination database version. The destination database version can be 4.0.x or later. The migration must be performed from an older database version to the same or a newer version. Check whether the features of the newer version are used in the source database. If the new features are not supported in the destination database, the synchronization may fail. |
| <b>Failure Cause and Handling Suggestion</b> | <p>Item to be confirmed: The versions of the source and destination databases are different.</p> <p>Potential problem: The source database version is 4.4.5 and the destination database version is 4.0.3, which does not comply with the migration version requirements. If the source database supports features of the newer version but the destination database does not, the synchronization may fail.</p>                                    |
|                                              | <p>Failure cause: The destination database version does not meet the requirements.</p> <p>Handling suggestion: Select version 4.0.x for the destination database because the source database version is later than 4.0.x and the destination database version is 3.4.x.</p>                                                                                                                                                                         |

### 6.1.7 Networks



### 6.1.7.1 Checking Whether the Source Database Is Connected

#### MySQL Migration

**Table 6-61** Checking whether the source database is connected

|                                              |                                                                                                                                                                                                                                                                                    |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database is connected                                                                                                                                                                                                                                           |
| <b>Description</b>                           | Check the connectivity and accuracy of the IP address, port number, username, and password of the source database.                                                                                                                                                                 |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The connection fails.<br>Handling suggestion: Configure the network by referring to "Network Types" in Real-Time Migration.                                                                                                                                         |
|                                              | Failure cause: Incorrect username or password<br>Handling suggestion: Check whether the input username and password for the connection test are correct.                                                                                                                           |
|                                              | Failure cause: The database account does not allow remote connections.<br>Handling suggestion:<br>Run the following command to create a user that allows remote connections. After the migration, delete this user.<br><b>CREATE USER 'Account' @ '%' IDENTIFIED BY 'Password'</b> |
|                                              | Failure cause: The SSL CA root certificate is invalid.<br>Handling suggestion: Upload a valid SSL CA certificate.                                                                                                                                                                  |
|                                              | Failure cause: No SSL CA root certificate exists.<br>Handling suggestion: Contact technical support.                                                                                                                                                                               |
|                                              | Failure cause: The database is unavailable.<br>Handling suggestion: Contact technical support.                                                                                                                                                                                     |

#### PostgreSQL Synchronization

**Table 6-62** Checking whether the source database is connected

|                                              |                                                                                                                                                      |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database is connected                                                                                                             |
| <b>Description</b>                           | Check the connectivity and accuracy of the IP address, port number, username, and password of the source database.                                   |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The IP address is inaccessible.<br>Handling suggestion: Configure the network by referring to "Network Types" in Real-Time Migration. |

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Handling Suggestion</b> | <p>Failure cause: The connection fails.</p> <p>Handling suggestion: Configure the network by referring to "Network Types" in Real-Time Migration.</p>                                                                                                                                                                                                                                                                                                                 |
|                            | <p>Failure cause: The database account does not allow remote connections.</p> <p>Handling suggestion:</p> <p>Configure the remote connection permission for the user in the <b>pg_hba.conf</b> file.</p> <p>Open <b>pg_hba.conf</b>, add the following, and restart the database for the modification to take effect:</p> <pre>host all xxx(dbuser) 0.0.0.0/0 method</pre> <p>After the migration is complete, delete this record and restart the database again.</p> |
|                            | <p>Failure cause: Failed to connect to the database.</p> <p>Handling suggestion:</p> <p>The <b>listen_addresses</b> parameter value or port number in the <b>postgres.conf</b> file is incorrect.</p> <p>In the <b>postgres.conf</b> file, set the <b>listen_addresses</b> value to '*' or set the port number to the correct value. Then, restart the database for the modification to take effect.</p>                                                              |
|                            | <p>Failure cause: Incorrect username or password</p> <p>Handling suggestion: Check whether the input username and password for the connection test are correct.</p>                                                                                                                                                                                                                                                                                                   |
|                            | <p>Failure cause: The user does not have the login permission.</p> <p>Handling suggestion:</p> <p>Run the following command to grant the login permission to the user:</p> <pre><b>alter role xxx(dbuser) login</b></pre>                                                                                                                                                                                                                                             |
|                            | <p>Failure cause: The postgres database does not exist in the source database.</p> <p>Handling suggestion: Create a postgres database.</p>                                                                                                                                                                                                                                                                                                                            |

## MongoDB Migration

**Table 6-63** Checking whether the source database is connected

|                    |                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>  | Whether the source database is connected                                                                           |
| <b>Description</b> | Check the connectivity and accuracy of the IP address, port number, username, and password of the source database. |

|                                              |                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The connection fails.<br>Handling suggestion: Configure the network by referring to "Network Types" in Real-Time Migration.                                                                                                                                                              |
|                                              | Failure cause: The username, password, or authentication database of the source database is incorrect.<br>Handling suggestion: Check that the input username, password, and authentication database for the connection test are correct.                                                                |
|                                              | Failure cause: The database is unavailable.<br>Handling suggestion: Contact technical support.                                                                                                                                                                                                          |
|                                              | Failure cause: SSL connection parameters are not configured.<br>Handling suggestion: Contact technical support.                                                                                                                                                                                         |
|                                              | Failure cause: The source database cannot connect to the port of the replication instance.<br>Handling suggestion: Modify the firewall, security group, or ACL configurations of the source and destination databases to enable the source database to connect to the port of the replication instance. |

## 6.1.7.2 Checking Whether the Destination Database Is Connected

### MySQL Migration

**Table 6-64** Checking whether the destination database is connected

|                                              |                                                                                                                                                                                                                                                                                    |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database is connected                                                                                                                                                                                                                                      |
| <b>Description</b>                           | Check the connectivity and accuracy of the IP address, port number, username, and password of the destination database.                                                                                                                                                            |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The connection fails.<br>Handling suggestion: Configure the network by referring to "Network Types" in section Real-Time Migration.                                                                                                                                 |
|                                              | Failure cause: Incorrect username or password<br>Handling suggestion: Check whether the input username and password for the connection test are correct.                                                                                                                           |
|                                              | Failure cause: The database account does not allow remote connections.<br>Handling suggestion:<br>Run the following command to create a user that allows remote connections. After the migration, delete this user.<br><b>CREATE USER 'Account' @ '%' IDENTIFIED BY 'Password'</b> |

|  |                                                                                                           |
|--|-----------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: The database is unavailable.</p> <p>Handling suggestion: Contact technical support.</p> |
|--|-----------------------------------------------------------------------------------------------------------|

## PostgreSQL Synchronization

**Table 6-65** Checking whether the destination database is connected

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database is connected                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>                           | Check the connectivity and accuracy of the IP address, port number, username, and password of the destination database.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The IP address is inaccessible.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                                              | Failure cause: The connection fails.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|                                              | <p>Failure cause: The database account does not allow remote connections.</p> <p>Handling suggestion:</p> <p>Grant the remote connection permission for the user in the <b>pg_hba.conf</b> file because the replication instance and user are not configured in the <b>pg_hba.conf</b> configuration file.</p> <p>Add the following to the <b>pg_hba.conf</b> configuration file. After the migration is complete, delete this record and restart the database for the modification to take effect.</p> <pre>host all xxx(dbuser) 0.0.0.0/0 method</pre> |
|                                              | <p>Failure cause: Failed to connect to the database. The failure may be caused by the incorrect <b>listen_addresses</b> parameter value or port number in <b>postgres.conf</b>.</p> <p>Handling suggestion: In the <b>postgres.conf</b> file, set <b>listen_addresses</b> to "*" or set the port number to the correct value. Then, restart the database for the modification to take effect.</p>                                                                                                                                                        |
|                                              | <p>Failure cause: Incorrect username or password</p> <p>Handling suggestion: Check whether the input username and password for the connection test are correct.</p>                                                                                                                                                                                                                                                                                                                                                                                      |
|                                              | <p>Failure cause: The user does not have the login permission.</p> <p>Handling suggestion:</p> <p>Run the following command to grant the login permission to the user:</p> <pre><b>alter role xxx(dbuser) login</b></pre>                                                                                                                                                                                                                                                                                                                                |
|                                              | <p>Failure cause: The postgres database does not exist in the source database.</p> <p>Handling suggestion: Create a postgres database.</p>                                                                                                                                                                                                                                                                                                                                                                                                               |

## MongoDB Migration

**Table 6-66** Checking whether the destination database is connected

|                                              |                                                                                                                                                          |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database is connected                                                                                                            |
| <b>Description</b>                           | Check the connectivity and accuracy of the IP address, port number, username, and password of the destination database.                                  |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The connection fails.                                                                                                                     |
|                                              | Failure cause: Incorrect username or password<br>Handling suggestion: Check whether the input username and password for the connection test are correct. |
|                                              | Failure cause: The database is unavailable.<br>Handling suggestion: Contact technical support.                                                           |
|                                              | Failure cause: SSL connection parameters are not configured.<br>Handling suggestion: Contact technical support.                                          |

### 6.1.7.3 Checking Whether the Destination Database Can Connect to the Source Database

## MySQL Migration and Synchronization

**Table 6-67** Checking whether the destination database can connect to the source database

|                                              |                                                                                                                                                                                        |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database can connect to the source database                                                                                                                    |
| <b>Description</b>                           | Check whether the destination database can connect to the source database.                                                                                                             |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The destination database fails to connect to the source database.<br>Handling suggestion: Configure the network by referring to "Network Types" in Real-Time Migration. |

## 6.1.8 Database Objects

### 6.1.8.1 Checking Whether the Source Database Contains a MyISAM Table

#### MySQL

**Table 6-68** Checking whether the source database contains a MyISAM table

|                                                     |                                                                                                                                                                                                                                                                                                         |
|-----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the source database contains a MyISAM table                                                                                                                                                                                                                                                     |
| <b>Description</b>                                  | If the source database contains a MyISAM table, the migration will fail.                                                                                                                                                                                                                                |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Item to be confirmed: The source database contains MyISAM tables that are not supported by the destination database, which may cause the migration to fail.<br>Handling suggestion: Convert the tables in the source database to InnoDB tables and try again. Alternatively, contact technical support. |

### 6.1.8.2 Checking Whether the Source Database Contains the Functions or Stored Procedures that the Source Database User Is Not Authorized to Migrate

#### MySQL

**Table 6-69** Checking whether the source database contains the functions or stored procedures that the source database user is not authorized to migrate

|                                              |                                                                                                                                                                                                            |
|----------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database contains the functions or stored procedures that the source database user is not authorized to migrate.                                                                        |
| <b>Description</b>                           | The source database contains the functions or stored procedures that the source database user is not authorized to migrate.                                                                                |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The source database user does not have the permission to migrate functions and stored procedures.<br>Handling suggestion: Ensure that the source database user has the highest-level right. |

### 6.1.8.3 Checking Whether the Source Database Tables Use Storage Engines Not Supported by the Destination Database

#### MySQL Migration

**Table 6-70** Checking whether the source database tables use storage engines not supported by the destination database

|                                                     |                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the source database tables use storage engines not supported by the destination database                                                                                                                                                                           |
| <b>Description</b>                                  | Check whether the source database tables use storage engines not supported by the destination database. If yes, the migration fails.                                                                                                                                       |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Failure cause: The source database tables use the storage engines that are not supported by the destination database.<br>Handling suggestion: Go back to the previous page and deselect the tables that use the storage engines not supported by the destination database. |

### 6.1.8.4 Checking Whether the Source Database Tables Contain Primary Keys

#### MySQL Migration

**Table 6-71** Checking whether the source database tables contain primary keys

|                                                     |                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the source database tables contain primary keys                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                  | If tables to be migrated in the source database do not contain primary keys, the migration may fail.                                                                                                                                                                                                                                                                       |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Item to be confirmed: The tables to be migrated in the source database do not contain primary keys.<br>Handling suggestion: Create a primary key for the table. If the table does not have a primary key to uniquely identify every row and the network connection is unstable, the data in the destination database may be inconsistent with that in the source database. |

## 6.1.8.5 Checking Whether the Source Database Contains Triggers or Events

### MySQL Migration

**Table 6-72** Checking whether the source database contains triggers or events

|                                                     |                                                                                                                                                                                                                                                                                                                   |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the source database contains triggers or events                                                                                                                                                                                                                                                           |
| <b>Description</b>                                  | To prevent unexpected operations on the destination database automatically triggered by triggers or events, this task starts the trigger or event migration only after you stop the task. If you close or disconnect the source database connection during the task running, triggers or events are not migrated. |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Item to be confirmed: The source database contains triggers or events.<br>Handling suggestion: Stop the task first and then disconnect the network to ensure the completeness of the migration.                                                                                                                   |

## 6.1.9 Database Configuration Items

### 6.1.9.1 Checking Whether the Source Database Name Is Valid

#### MySQL Migration

**Table 6-73** Checking whether the source database name is valid

|                                              |                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database name is valid                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                           | The source database name cannot contain invalid characters. It must contain 1 to 64 characters, including only lowercase letters, digits, hyphens (-), and underscores (_).<br>If the source database name contains any invalid character, the migration fails.                                                                                                                              |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                                                                                                                                                                                        |
|                                              | Failure cause: The source database name cannot contain invalid characters. It must contain 1 to 64 characters, including only lowercase letters, digits, hyphens (-), and underscores (_).<br>Handling suggestion: Change the source database names that contain unsupported characters or go back to the previous page and select the databases that do not contain unsupported characters. |



### 6.1.9.2 Checking Whether the Source Database Table Name Is Valid

#### MySQL Migration

**Table 6-74** Checking whether the source database table name is valid

|                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the source database table name is valid                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                           | If the source database table name contains invalid character, the synchronization task fails.                                                                                                                                                                                                                                                                                                                                    |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The source database table names contain unsupported characters, non-ASCII characters, or the following characters: ></\<br>Handling suggestion: To solve this problem, perform the following steps:<br>Click <b>Previous</b> to return to the <b>Select Migration Type</b> page. Select a customized object and do not select the table that contains unsupported characters.<br>Method 2: Change the table name. |

### 6.1.9.3 Checking Whether the Source Database View Name Is Valid

#### MySQL

**Table 6-75** Checking whether the source database contains view names with non-ASCII characters

|                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the source database contains view names with non-ASCII characters                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>                                  | If the source database contains non-ASCII characters, the migration will fail.                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Failure cause: The source database view names contain unsupported characters, non-ASCII characters, or the following characters: ></\<br>Handling suggestion: To solve this problem, perform the following steps:<br>Method 1:<br>Click <b>Previous</b> to return to the <b>Select Migration Type</b> page. Select a customized object and do not select the view name that contains unsupported characters.<br>Method 2: Change the view name. |

## 6.1.9.4 Checking Whether the Shard Key Can Be Obtained from the Source Database

### MongoDB Migration

**Table 6-76** Checking whether the shard keys can be obtained from the source database

|                                              |                                                                                                                                                        |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the shard keys can be obtained from the source database                                                                                        |
| <b>Description</b>                           | Check whether the destination database user permissions meet the migration requirements. If the permissions are insufficient, the migration will fail. |
| <b>Failure Cause and Handling Suggestion</b> | Item to be confirmed: The source database is a replica set but the shard keys have not been configured in the destination database.                    |
|                                              | Item to be confirmed: The source database type is unknown, and the shard keys have not been configured in the destination database.                    |
|                                              | Item to be confirmed: The source database contains collections that do not have shard keys configured.                                                 |

## 6.1.10 Conflicts

### 6.1.10.1 Checking Whether the Names of the Source and Destination Databases Are the Same

#### MySQL Migration

**Table 6-77** Checking whether the names of the source and destination databases are the same

|                                              |                                                                                                                                                                                 |
|----------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the names of the source and destination databases are the same                                                                                                          |
| <b>Description</b>                           | Check whether the names of the source and destination databases are the same.                                                                                                   |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.           |
|                                              | Failure cause: This item cannot be checked because the destination database fails to be connected.<br>Handling suggestion: Check whether the destination database is connected. |

|  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|--|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Failure cause: Insufficient user permissions</p> <p>Handling suggestion: Check whether the database user permissions meet the migration requirements.</p>                                                                                                                                                                                                                                                                                                              |
|  | <p>Handling suggestion:</p> <ul style="list-style-type: none"> <li>• If you are migrating data to the cloud, determine whether to delete the databases with the same names as the source databases or specify a new destination DB instance based on site requirements.</li> <li>• If you are migrating data out of the cloud, determine whether to use the original destination database or specify a new destination DB instance based on site requirements.</li> </ul> |
|  | <p>Failure cause: During an incremental migration, the source and destination databases cannot have the same names.</p> <p>Handling suggestion: Determine whether to retain these databases in the destination RDS DB instance or specify another destination RDS DB instance.</p>                                                                                                                                                                                        |

### 6.1.10.2 Checking Whether the Destination Database Contains a Non-Empty Collection with the Same Name As the Source Database

#### MongoDB Migration

**Table 6-78** Checking whether the destination database contains a non-empty collection with the same name as the source database

|                                              |                                                                                                                                                                                                                                                                            |
|----------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the destination database contains a non-empty collection with the same name as the source database                                                                                                                                                                 |
| <b>Description</b>                           | Check whether the source and destination databases have the same non-empty collections to prevent existing databases from being overwritten. If the source and destination databases have the same collections, the migration cannot be performed.                         |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: This item cannot be checked because the source database failed to be connected.</p> <p>Handling suggestion: Check whether the source database is connected.</p>                                                                                          |
|                                              | <p>Failure cause: This item cannot be checked because the destination database failed to be connected.</p> <p>Handling suggestion: Check whether the destination database is connected.</p>                                                                                |
|                                              | <p>Failure cause: The same non-empty collections exist in both the source and destination databases.</p> <p>Handling suggestion: Determine whether to delete the same non-empty collections from the destination DB instance or specify a new destination DB instance.</p> |

## 6.1.11 SSL Connections

### 6.1.11.1 Checking Whether the SSL Connection Is Correctly Configured

#### MySQL

**Table 6-79** Checking whether the SSL connection is correctly configured

|                                              |                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the SSL connection is correctly configured                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>                           | Check whether the SSL connection is correctly configured for the source database.                                                                                                                                                                                                                                                                   |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: This item cannot be checked because the source database fails to be connected.<br>Handling suggestion: Check whether the source database is connected.                                                                                                                                                                               |
|                                              | Failure cause: Insufficient user permissions<br>Handling suggestion: Check whether the database user permissions meet the migration requirements.                                                                                                                                                                                                   |
|                                              | Failure cause: The database is unavailable.<br>Handling suggestion: Contact technical support.                                                                                                                                                                                                                                                      |
|                                              | Item to be confirmed: The source database user must have the REQUIRE SSL permission when using the SSL connection.<br>Handling suggestion: This alarm does not affect the migration process. If you require the SSL connection, you are advised to grant the REQUIRE SSL permission to the source database user.                                    |
|                                              | Item to be confirmed: The destination database user must have the REQUIRE SSL permission when using the SSL connection.<br>Handling suggestion: This alarm does not affect the migration process. If you require the SSL connection, you are advised to grant the REQUIRE SSL permission to the destination database user.                          |
|                                              | Failure cause: The source database user has the REQUIRE SSL permission but did not upload the encryption certificate. The SSL connection cannot be used.<br>Handling suggestion: On the <b>Configure Source and Destination Databases</b> page, enable the SSL connection and upload the certificate, or change the source database user.           |
|                                              | Failure cause: The destination database user has the REQUIRE SSL permission but did not upload the encryption certificate. The SSL connection cannot be used.<br>Handling suggestion: On the <b>Configure Source and Destination Databases</b> page, enable the SSL connection and upload the certificate, or change the destination database user. |

|  |                                                                                                                                                                                                                                                                                                                                                                                                |
|--|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  | <p>Item to be confirmed: Currently, the SSL connection is not enabled. DRS must ensure that the source database account allows non-SSL connections to the source database.</p> <p>Handling suggestion: Manually check whether the source database account allows non-SSL connections, or try to perform a migration. (By default, the source database account allows non-SSL connections.)</p> |
|  | <p>Failure cause: The SSL connection is enabled for the source database but no certificate has been uploaded.</p> <p>Handling suggestion: On the <b>Configure Source and Destination Databases</b> page, upload a certificate or disable the SSL connection for the source database.</p>                                                                                                       |

### 6.1.11.2 Checking Whether the SSL Connection Is Enabled for the Source Database

#### PostgreSQL

**Table 6-80** Checking whether the SSL connection is enabled for the source database

|                                              |                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                            | Whether the SSL connection is enabled for the source database                                                                                                                                                                                                                                                                              |
| <b>Description</b>                           | Check whether the SSL connection is enabled for the source database.                                                                                                                                                                                                                                                                       |
| <b>Failure Cause and Handling Suggestion</b> | <p>Failure cause: The source database SSL connection is disabled.</p> <p>Handling suggestion: In the <b>postgresql.conf</b> file, set <b>ssl_ca_file</b> to the directory of an SSL root CA certificate and set <b>ssl</b> to <b>on</b> to enable the SSL connection. Then, restart the database for the modifications to take effect.</p> |

### 6.1.11.3 Checking Whether the SSL Certificate of the Destination Database Exists

#### MySQL

**Table 6-81** Checking whether the SSL certificate of the destination database exists

|                   |                                                                |
|-------------------|----------------------------------------------------------------|
| <b>Check Item</b> | Whether the SSL certificate of the destination database exists |
|-------------------|----------------------------------------------------------------|

|                                              |                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>                           | Check whether the SSL certificate type of the destination database is correct during migration. Otherwise, the migration fails.                                                                                                                                                                                                                                             |
| <b>Failure Cause and Handling Suggestion</b> | Failure cause: The SSL certificate of the destination database does not exist.<br>Handling suggestion: On the <b>Configure Source and Destination Databases</b> page, enable SSL connection for the destination database and upload an encryption certificate that contains only one beginning tag <b>BEGIN CERTIFICATE</b> and one end tag <b>END CERTIFICATE</b> .        |
|                                              | Failure cause: The SSL certificate type of the destination database is not supported.<br>Handling suggestion: On the <b>Configure Source and Destination Databases</b> page, enable SSL connection for the destination database and upload an encryption certificate that contains only one beginning tag <b>BEGIN CERTIFICATE</b> and one end tag <b>END CERTIFICATE</b> . |

## 6.1.12 Object Dependencies

### 6.1.12.1 Checking Whether Referenced Tables Are Selected for Migration

#### MySQL Migration

**Table 6-82** Checking whether the tables referenced by the foreign key in the table to be migrated are selected for migration.

|                                                     |                                                                                                                                                                       |
|-----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Check Item</b>                                   | Whether the tables referenced by the foreign key in the table to be migrated are selected for migration.                                                              |
| <b>Description</b>                                  | The tables referenced by the foreign key in the table to be migrated are not selected for migration.                                                                  |
| <b>Item to Be Confirmed and Handling Suggestion</b> | Failure cause: Tables referenced by the foreign key in the table to be migrated are not selected for migration.<br>Handling suggestion: Select the referenced tables. |

## 6.2 Failure Cases

## 6.2.1 Backup Migration Failed Because Backup Files Cannot Be Found

### Scenarios

When you migrate full backups from self-built OBS buckets to clouds, the following error message is displayed: restore:null.

### Possible Causes

The possible causes are as follows:

- Backup files are deleted after you submit a backup migration task.
- When you upload backup files to a self-built OBS bucket, you select **Archive** for **Storage Class**. OBS archive storage offers cloud storage for rarely accessed data. An archive file uploaded for the first time is in the **Not restored** status. As a result, a Microsoft SQL Server DB instance cannot download the file.

### Solutions

Based on the previous analysis, solutions are provided as follows:

#### Solution 1

If the migration fails because you delete the backup files, you can upload the deleted backup files again to a self-built OBS bucket and select **Standard** for **Storage Class**. For details, see **Uploading a File** in *Object Storage Service Console Operation Guide*.

#### Solution 2

- If the migration failed because the storage class of your backup files is **Archive**, perform the following steps. If the size of backup files is small, upload the backup files again to an OBS bucket and select **Standard** for **Storage Class**.
- If the backup files are large in size, log in to the OBS console and click the bucket to which the backup files are uploaded. On the displayed page, choose **Objects** in the navigation pane on the left. On the **Objects** page, select the object to be restored and click **Restore** above the file list. After the status of the backup files becomes **Restored**, submit an offline migration task again.

## 6.2.2 Backup Migration Failed Because a Backup Database Cannot Be Found in the Backup Files

### Scenarios

When you migrate full backups from self-built OBS buckets to clouds, the system displays an error message indicating that the migration failed because the source database cannot be found in the backup files.

## Possible Cause

The name of a .bak backup file uploaded to a self-built OBS bucket is too long.

## Solution

Based on the previous analysis, a solution is provided as follows:

- Step 1** Change the name of the backup file in the local database and upload the file to a self-built OBS bucket again.

----End



# A Change History

| Released On | Description                                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 2022-05-30  | This issue is the sixth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Supported real-time synchronization from PostgreSQL to PostgreSQL.</li></ul>                                                                                                                                                                                                        |
| 2021-11-30  | This issue is the fifth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Supported real-time migration from MongoDB to DDS.</li><li>• Supported real-time migration from DDS to MongoDB.</li></ul>                                                                                                                                                           |
| 2021-08-30  | This issue is the fourth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Supported real-time synchronization from MySQL to MySQL.</li><li>• Supported real-time synchronization from MySQL to Kafka.</li><li>• Supported real-time synchronization from Oracle to MySQL.</li><li>• Supported real-time synchronization from Oracle to PostgreSQL.</li></ul> |
| 2021-05-31  | This issue is the third official release, which incorporates the following change: <ul style="list-style-type: none"><li>• Supported backup migrations of Microsoft SQL Server databases.</li></ul>                                                                                                                                                                                                             |
| 2019-12-30  | This issue is the first official release.                                                                                                                                                                                                                                                                                                                                                                       |