

Data Replication Service

Service Overview

Issue 01
Date 2026-01-23



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2026. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

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

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

Contents

1 Infographics.....	1
1.1 DRS Overview.....	1
1.2 Five Major Functions.....	1
1.3 Data Comparison.....	2
1.4 User Permission Migration.....	2
2 What Is DRS?.....	3
3 Supported Databases.....	8
3.1 Real-Time Migration.....	8
3.2 Backup Migration.....	14
3.3 Real-Time Synchronization.....	16
3.4 Data Subscription.....	35
3.5 Real-Time Disaster Recovery.....	36
3.6 Workload Replay.....	37
4 Product Architecture and Function Principles.....	39
5 Advantages.....	44
6 Functions and Features.....	45
6.1 Real-Time Migration.....	45
6.2 Backup Migration.....	50
6.3 Real-Time Synchronization.....	54
6.4 Data Subscription.....	70
6.5 Real-Time Disaster Recovery.....	71
6.6 Workload Replay.....	73
7 Specification Description.....	81
7.1 Real-Time Synchronization.....	81
7.2 Real-Time Disaster Recovery.....	90
8 Mapping Data Types.....	93
8.1 MySQL->PostgreSQL.....	93
8.2 MySQL->GaussDB.....	95
8.3 DDM->GaussDB.....	99
8.4 MySQL->Oracle.....	102

8.5 Oracle->MySQL.....	103
8.6 Oracle->TaurusDB.....	105
8.7 Oracle->GaussDB.....	108
8.8 Oracle->DDM.....	119
8.9 Oracle->PostgreSQL.....	121
8.10 TaurusDB->Oracle.....	123
8.11 GaussDB->MySQL.....	124
8.12 GaussDB->PostgreSQL.....	127
8.13 GaussDB->Oracle.....	130
8.14 DB2 for LUW->GaussDB.....	131
8.15 PostgreSQL->GaussDB.....	132
8.16 TiDB->TaurusDB.....	134
8.17 Microsoft SQL Server->GaussDB.....	135
8.18 Microsoft SQL Server->MySQL.....	137
8.19 Microsoft SQL Server->TaurusDB.....	140
8.20 Microsoft SQL Server->PostgreSQL.....	143
9 Security.....	146
9.1 Shared Responsibilities.....	146
9.2 Identity Authentication and Access Control.....	148
9.3 Data Protection.....	148
9.4 Audit and Logs.....	149
9.5 Risk Monitoring.....	149
9.6 Fault Recovery.....	150
9.7 Certificates.....	150
9.8 Compliance Description.....	152
10 Permissions Management.....	153
11 IAM-based Permissions Management.....	168
12 Agency Management.....	178
13 Notes and Constraints.....	183
13.1 Constraints on Migration Tasks.....	183
13.2 Constraints on Synchronization Tasks.....	184
13.3 Constraints on DR Tasks.....	187
14 Accessing DRS.....	188
15 Related Services.....	189
16 Basic Concepts.....	192

1 Infographics

1.1 DRS Overview



1.2 Five Major Functions



1.3 Data Comparison



1.4 User Permission Migration



2 What Is DRS?

Data Replication Service (DRS) is an easy-to-use, stable, and efficient cloud service for online database migration and real-time database synchronization.

It simplifies data transfers between databases and reduces data transfer costs.

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

DRS provides multiple functions, including real-time migration, backup migration, real-time synchronization, data subscription, and real-time disaster recovery.

Introduction Video

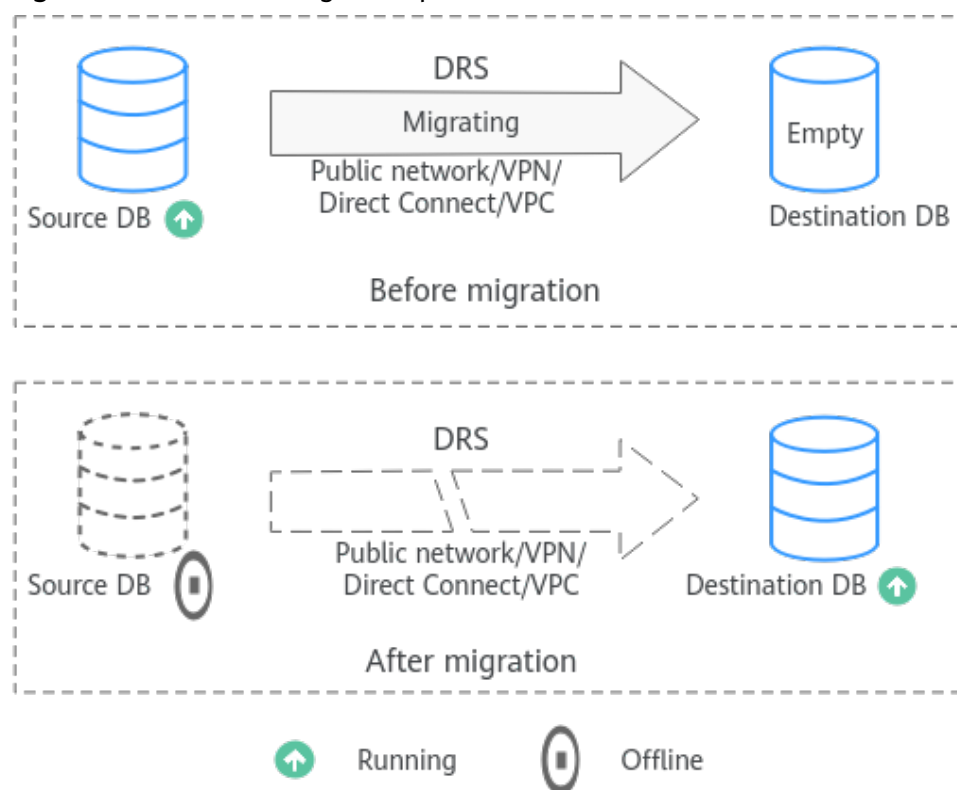
Real-Time Migration

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.

Real-time migration can be performed over different networks, such as public networks, Virtual Private Cloud (VPCs), Virtual Private Network (VPNs), and Direct Connect. With these network connections, you can migrate data 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.

Figure 2-1 Real-time migration process




Backup Migration

For security reasons, it is often necessary to hide the real IP address of your database. Migrating data through dedicated connections is an option, but it is expensive. DRS supports backup migration, which allows you to export data from your source database for backup and upload the backup files to Object Storage Service (OBS). Then, you can restore the backup files to the destination database to complete the migration. This method lets you migrate data without exposing your source databases to the Internet.

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.

Figure 2-2 Backup migration process



Real-Time Synchronization

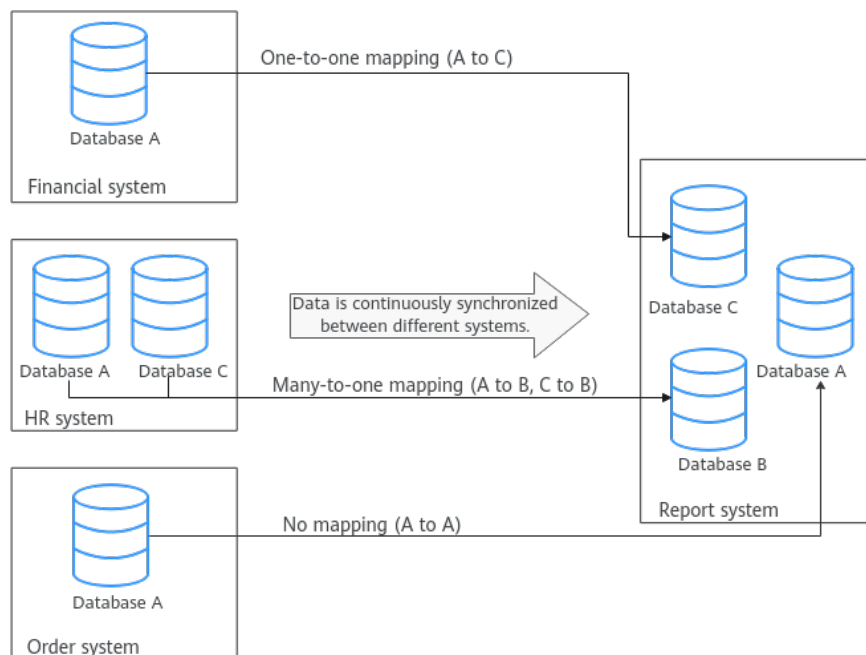
In a real-time synchronization, data of critical workloads can be moved from a source database to a destination database while consistency of data is guaranteed.

It is different from migration. Migration means moving your entire database from one platform to another, whereas, synchronization refers to the continuous flow of data between different workloads.

You can use real-time synchronization for real-time analysis and report systems.

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.

Figure 2-3 Many-to-one real-time synchronization process




Data Subscription

With DRS, you can subscribe changes made to key workloads in a database for downstream services to consume. DRS caches the changes and uses a unified SDK API to facilitate downstream services to subscribe to, obtain, and consume the changes, decoupling databases from downstream systems.

Data subscription can be used by Kafka to subscribe to MySQL incremental data.

Figure 2-4 Data subscription

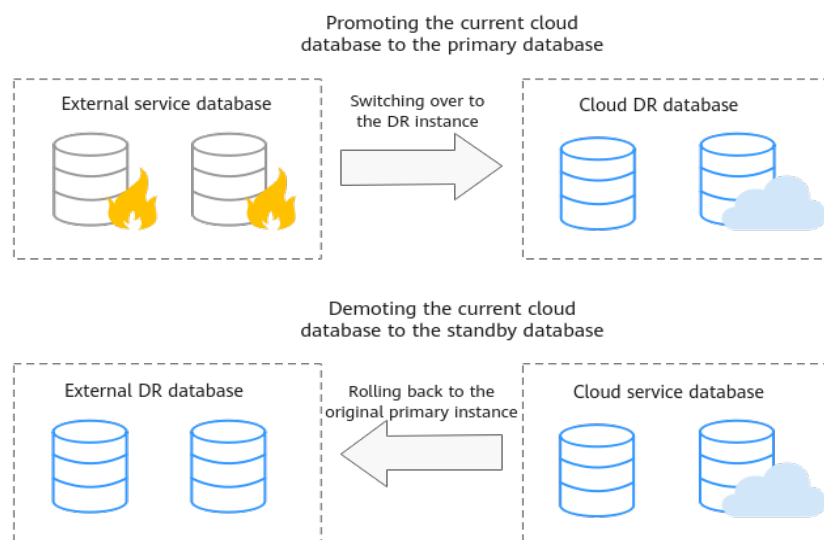


Real-Time Disaster Recovery

To prevent service unavailability caused by regional faults, DRS provides disaster recovery to ensure service continuity. You can easily implement disaster recovery between on-premises and cloud, without the need to invest a lot in infrastructure in advance.

The disaster recovery architectures, such as two-site three-data-center and two-site four-data center, are supported. A primary/standby switchover can be implemented by promoting a standby node or demoting a primary node in the disaster recovery scenario.

Figure 2-5 Real-time DR switchover



Workload Replay


A workload replay task simulates the source database's workload and runs the workload on the destination database, so you can evaluate the functionality and performance of the destination database.

A task consists of SQL capture and replay. All of the SQL statements (create, delete, update, and query operations) executed in the required period on the source database will be downloaded by a capture tool from the binlog, and then cached and injected into the destination database where you can trigger a replay and review performance.

Typical Scenarios

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

Figure 2-6 Workload replay



3 Supported Databases

DRS supports data flows between different DB engines. This section lists the supported databases in real-time migration, backup migration, real-time synchronization, data subscription, and real-time DR.

3.1 Real-Time Migration

You can migrate all database objects across cloud platforms, from on-premises databases to the cloud, or across regions on the cloud in real time. The following table lists the supported databases, versions, and migration types. For more information about real-time migration, see [Real-Time Migration](#).

Self-managed databases (such as MySQL and MongoDB) include on-premises databases and databases created on ECSs. RDS for MySQL refers to the MySQL databases created using Huawei Cloud RDS.

NOTE

- Some functions are in restricted use. To use them, [submit a service ticket](#).
- The destination database version must be the same as or later than the source database version.

MySQL Serving as the Source in Migration

Table 3-1 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
MySQL	<ul style="list-style-type: none"> Self-managed MySQL Versions 5.5, 5.6, 5.7, and 8.0 MySQL on other clouds Versions 5.5, 5.6, 5.7, and 8.0 	RDS for MySQL All versions	Full Full+Incremental
		DDM The same version as that of the associated RDS instance.	Full Full+Incremental
		TaurusDB Compatible with MySQL 8.0	Full Full+Incremental
	RDS for MySQL All versions	RDS for MySQL All versions	Full Full+Incremental
		Self-managed or other cloud MySQL Versions 5.5, 5.6, 5.7, and 8.0	Full Full+Incremental
		DDM The same version as that of the associated RDS instance.	Full Full+Incremental
		TaurusDB Compatible with MySQL 8.0	Full Full+Incremental

MongoDB Serving as the Source in Migration

Table 3-2 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
MongoDB	<ul style="list-style-type: none"> Self-managed MongoDB Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 MongoDB on other clouds Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 DDS Versions 3.2, 3.4, 4.0, 4.2, 4.4, and 5.0 <p>NOTE</p> <ul style="list-style-type: none"> The source does not support GeminiDB Mongo. 	<p>DDS</p> <p>Versions 3.4, 4.0, 4.2, 4.4, and 5.0</p>	<p>Full</p> <p>Full+Incremental migration supports the following scenarios:</p> <ul style="list-style-type: none"> Replica set -> Replica set Replica set -> Cluster Cluster -> Cluster Cluster -> Replica set <p>NOTE</p> <ul style="list-style-type: none"> If the source is a DDS cluster instance, the incremental migration can only be performed when the source and destination databases are in the same VPC. If the source database is a DDS 3.2 cluster instance, only full migration is supported. Incremental migration is not supported because a DDS 3.2 cluster instance does not have a shard IP address. To perform a full +incremental migration for a single node instance, the source database must be a Huawei Cloud single node instance.

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
MongoDB	<ul style="list-style-type: none"> Self-managed MongoDB Versions 3.2, 3.4, and 4.0 MongoDB on other clouds Versions 3.2, 3.4, and 4.0 DDS Versions 3.4 and 4.0 <p>NOTE</p> <ul style="list-style-type: none"> The source does not support GeminiDB Mongo. 	GeminiDB Mongo Versions 3.4 and 4.0	<p>Full</p> <p>Full+Incremental migration supports the following scenarios:</p> <ul style="list-style-type: none"> Replica set -> Replica set Replica set -> Cluster Cluster -> Cluster <p>NOTE</p> <p>If the source is a DDS cluster, only full migration is supported.</p>

AWS DocumentDB Serving as the Source in Migration

Table 3-3 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
AWS Document	Replica set 4.0, replica set 5.0, and cluster 5.0	DDS instances (versions 3.4, 4.0, 4.2, 4.4, and 5.0)	<ul style="list-style-type: none"> Full <p>Full+Incremental migration supports the following scenarios:</p> <ul style="list-style-type: none"> Replica set -> Replica set Replica set -> Cluster Cluster -> Cluster Cluster -> Replica set <p>NOTE</p> <ol style="list-style-type: none"> If the source database is a replica set, full+incremental, full, and incremental migration tasks are supported. If the source database is a cluster, only full migration tasks are supported. Currently, only synchronization tasks with AWS DocumentDB serving as the source are supported.

DDS Serving as the Source in Migration

Table 3-4 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
DDS	DDS Versions 3.2, 3.4, 4.0, 4.2, 4.4, and 5.0	<ul style="list-style-type: none"> Self-managed MongoDB Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 MongoDB on other clouds Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 	<p>Full Full+Incremental migration supports the following scenarios:</p> <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> <ol style="list-style-type: none"> If the source database is a replica set, full +incremental and full migration tasks are supported. If the source database is a cluster, only full migration tasks are supported.

MySQL Schema and Logic Table Serving as the Source in Migration

Table 3-5 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
MySQL schema and logic table	DDM	DDM	Full Full+Incremental

Redis Serving as the Source in Migration

Table 3-6 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
Single-Node or Primary/Standby Redis	Self-managed official open-source single-node or primary/standby Redis Versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.x	GeminiDB Redis	Full Full+Incremental

Redis Cluster Serving as the Source in Migration

Table 3-7 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
Redis Cluster	<ul style="list-style-type: none"> On-premises open-source Redis Cluster Versions 3.0.x, 4.0.x, and 5.0.x ECS-hosted open-source Redis Cluster Versions 3.0.x, 4.0.x, and 5.0.x 	GeminiDB Redis	Full+Incremental

GeminiDB Redis Serving as the Source in Migration

Table 3-8 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Migration Type
GeminiDB Redis	GeminiDB Redis	<ul style="list-style-type: none">Self-managed open-source Codis Version 3.0 or laterSelf-managed open-source single-node Redis Versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, and 5.0.xSelf-managed open-source primary/standby Redis Versions 4.0.x and 5.0.x	Full+Incremental
	GeminiDB Redis	<ul style="list-style-type: none">Self-managed open-source Redis Cluster Version 4.0 or later	Full+Incremental

3.2 Backup Migration

You can export data from the source database to a backup file, upload the backup file to OBS, and then restore the backup file to the destination database. In this way, data migration can be complete without exposing your source database to the Internet. [Table 3-9](#) lists the supported databases, versions, and migration types. For more information about backup migration, see [Backup Migration](#).

NOTE

The major version of the destination database must be the same as or later than that of the source database.

Table 3-9 Database information

Backup File Version	Destination DB Version	Migration Method	Backup File Source
<p>On-premises and cloud Microsoft SQL Server backup file versions:</p> <ul style="list-style-type: none"> • Microsoft SQL Server 2000 Enterprise Edition and Standard Edition • Microsoft SQL Server 2005 Enterprise Edition and Standard Edition • Microsoft SQL Server 2008 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2012 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2014 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2016 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2017 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2019 Enterprise Edition, Standard Edition, and Web Edition 	<p>RDS for SQL Server</p> <ul style="list-style-type: none"> • Microsoft SQL Server 2008 Enterprise Edition, Standard Edition, and Web Edition (Existing versions) • Microsoft SQL Server 2012 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2014 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2016 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2017 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2019 Enterprise Edition, Standard Edition, and Web Edition 	<p>Full Incremental</p>	<p>OBS bucket RDS full backup</p>

Backup File Version	Destination DB Version	Migration Method	Backup File Source
RDS for SQL Server full backup file versions: <ul style="list-style-type: none">• Microsoft SQL Server 2008 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2012 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2014 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2016 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2017 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2019 Enterprise Edition, Standard Edition, and Web Edition		Full Incremental	OBS bucket RDS full backup

3.3 Real-Time Synchronization

Real-time synchronization refers to the process of copying data from one data source to another database while keeping data consistent. In this way, data of critical workloads can move between different systems in real time. The following table lists the supported databases, versions, and synchronization types. For more information about real-time synchronization, see [Real-Time Synchronization](#).

Self-managed databases (such as MySQL, Oracle, and PostgreSQL) include on-premises databases and databases created on ECSs. RDS for MySQL refers to the MySQL databases created using Huawei Cloud RDS.

NOTE

- Some functions are in restricted use. To use them, [submit a service ticket](#).
- The destination database version must be the same as or later than the source database version.

MySQL Serving as the Source in One-Way Synchronization

Table 3-10 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
MySQL	<ul style="list-style-type: none"> Self-managed MySQL Versions 5.5, 5.6, 5.7, and 8.0 MySQL on other clouds Versions 5.5, 5.6, 5.7, and 8.0 	RDS for MySQL All versions	Incremental Full Full+Incremental
		TaurusDB Compatible with MySQL 8.0	Incremental Full+Incremental
		RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Full Full+Incremental
		GaussDB Distributed (MySQL-compatible) The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions	Incremental Full Full+Incremental
		GaussDB Centralized (B-compatible or M-compatible) The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions	Incremental Full Full+Incremental
		Kafka Version 0.11 or later	Incremental Full+Incremental
		RDS for MariaDB Version 10.5	Incremental Full Full+Incremental
	RDS for MySQL All versions	RDS for MySQL All versions	Incremental Full+Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
		<ul style="list-style-type: none"> Self-managed MySQL Versions 5.5, 5.6, 5.7, and 8.0 MySQL on other clouds Versions 5.5, 5.6, 5.7, and 8.0 	Incremental Full+Incremental
		TaurusDB Compatible with MySQL 8.0	Incremental Full+Incremental
		RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Full Full+Incremental
		GaussDB Distributed (MySQL-compatible) The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions	Incremental Full Full+Incremental
		GaussDB Centralized (B-compatible or M-compatible) The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions	Incremental Full Full+Incremental
		Kafka Version 0.11 or later	Incremental Full+Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
		<ul style="list-style-type: none"> Self-managed MariaDB Version 10.0, 10.1, 10.2, 10.3, 10.4, or 10.5 MariaDB on other clouds Version 10.0, 10.1, 10.2, 10.3, 10.4, or 10.5 	Incremental Full Full+Incremental
		Self-managed Oracle Versions 10g, 11g, 12c, 18c, and 19c	Full+Incremental

MySQL Serving as the Source in Two-Way Synchronization

Table 3-11 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
MySQL	<ul style="list-style-type: none"> Self-managed MySQL Versions 5.5, 5.6, 5.7, and 8.0 MySQL on other clouds Versions 5.5, 5.6, 5.7, and 8.0 RDS for MySQL All versions 	RDS for MySQL All versions	Full+Incremental NOTE Two-way synchronization is in restricted use. To use it, submit a service ticket .

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
	RDS for MySQL All versions	<ul style="list-style-type: none"> Self-managed MySQL Versions 5.5, 5.6, 5.7, and 8.0 MySQL on other clouds Versions 5.5, 5.6, 5.7, and 8.0 RDS for MySQL All versions 	Full+Incremental NOTE Two-way synchronization is in restricted use. To use it, submit a service ticket .

MariaDB Serving as the Source in Synchronization

Table 3-12 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
MariaDB	<ul style="list-style-type: none"> On-premises MariaDB 10.3, 10.4, and 10.5 ECS-hosted MariaDB 10.3, 10.4, and 10.5 Other cloud MariaDB 10.3, 10.4, and 10.5 	RDS for MariaDB Version 10.5	Full+Incremental
	RDS for MariaDB Version 10.5	<ul style="list-style-type: none"> On-premises MariaDB Version 10.5 MariaDB built on ECSs Version 10.5 MariaDB on other clouds Version 10.5 	Full+Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
MariaDB	<ul style="list-style-type: none"> ● On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ● ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ● MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 on other clouds 	RDS for MySQL Versions 5.6, 5.7, and 8.0	Incremental Full+Incremental
MariaDB	<ul style="list-style-type: none"> ● On-premises MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ● ECS-hosted MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 ● MariaDB 10.0, 10.1, 10.2, 10.3, 10.4 and 10.5 on other clouds 	TaurusDB Compatible with MySQL 8.0	Incremental Full+Incremental

PostgreSQL Serving as the Source in Synchronization

Table 3-13 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
PostgreSQL	<ul style="list-style-type: none"> Self-managed PostgreSQL Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 PostgreSQL on other clouds Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 	RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Incremental Full Full +Incremental
	<ul style="list-style-type: none"> Self-managed PostgreSQL Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, 14, and 16 PostgreSQL on other clouds Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, 14, and 16 RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, and 16 	GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Incremental Full Full +Incremental NOTE Incremental synchronization is in restricted use. To use it, submit a service ticket.
		GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Incremental Full Full +Incremental NOTE Incremental synchronization is in restricted use. To use it, submit a service ticket.
	RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Kafka Version 0.11 or later	Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
	RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Self-managed PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Incremental Full Full+Incremental
	<ul style="list-style-type: none"> Self-managed PostgreSQL Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 PostgreSQL on other clouds Versions 9.4, 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 	Kafka Version 0.11 or later	Incremental

Oracle Serving as the Source in Synchronization

Table 3-14 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
Oracle	Self-managed Oracle Versions 10g, 11g, 12c, 18c, 19c, and 21c	RDS for MySQL All versions	Incremental Full Full+Incremental
		TaurusDB Compatible with MySQL 8.0	Full Full+Incremental
		RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16	Full Full+Incremental
		GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Incremental Full Full+Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
		GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Incremental Full Full+Incremental
		DDM	Full Full+Incremental
		Kafka Version 0.11 or later	Incremental

DDM Serving as the Source in Synchronization

Table 3-15 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
DDM	DDM	RDS for MySQL All versions	Incremental Full Full+Incremental
		Self-managed or other cloud MySQL Versions 5.6, 5.7, and 8.0	Incremental Full Full+Incremental
		DDM Based on the live network	Full+Incremental
		GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions	Full Full+Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
		GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions	Full Full+Incremental
		Self-managed Oracle Versions 10g, 11g, 12c, 18c, and 19c	Incremental Full Full+Incremental
		Kafka Version 0.11 or later	Incremental

TaurusDB Serving as the Source in Synchronization

Table 3-16 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
TaurusDB	TaurusDB Version 8.0	<ul style="list-style-type: none"> RDS for MySQL Version 8.0 Self-managed or other cloud MySQL Version 8.0 	Incremental Full+Incremental
		TaurusDB Compatible with MySQL 8.0	Incremental Full+Incremental
		Kafka Version 0.11 or later	Incremental Full+Incremental
		Self-managed Oracle Versions 10g, 11g, 12c, 18c, and 19c	Full+Incremental

GaussDB Distributed Serving as the Source in Synchronization

Table 3-17 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
TaurusDB Distributed	GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	RDS for MySQL Versions 5.6, 5.7, and 8.0	Full Incremental Full+Incremental
		Self-managed or other cloud MySQL Versions 5.5, 5.6, 5.7, and 8.0	Full Incremental Full+Incremental
		Self-managed Oracle Versions 10g, 11g, 12c, 18c, and 19c	Full Incremental Full+Incremental
		Kafka Version 0.11 or later	Incremental
		GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Incremental Full+Incremental
		GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Incremental Full+Incremental
		RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 Self-managed PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 NOTE This function is in restricted use. To use it, submit a service ticket .	Full Incremental Full+Incremental

GaussDB Centralized Serving as the Source in Synchronization

Table 3-18 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
GaussDB Centralized	GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	RDS for MySQL Versions 5.6, 5.7, and 8.0	Full Incremental Full+Incremental
		Self-managed or other cloud MySQL Versions 5.5, 5.6, 5.7, and 8.0	Full Incremental Full+Incremental
		Self-managed Oracle Versions 10g, 11g, 12c, 18c, and 19c	Full Incremental Full+Incremental
		Kafka Version 0.11 or later	Incremental
		GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Incremental Full+Incremental
		GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Incremental Full+Incremental
		RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 Self-managed PostgreSQL 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 NOTE This function is in restricted use. To use it, submit a service ticket.	Full Incremental Full+Incremental

MongoDB Serving as the Source in Synchronization

Table 3-19 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
MongoDB	<ul style="list-style-type: none"> Self-managed MongoDB Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 MongoDB on other clouds Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 DDS Versions 3.2, 3.4, 4.0, 4.2, 4.4, and 5.0 <p>NOTE The source does not support GeminiDB Mongo.</p>	DDS Versions 3.4, 4.0, 4.2, 4.4, and 5.0	Full+Incremental synchronization supports the following scenarios: <ul style="list-style-type: none"> Replica set -> Replica set Replica set -> Cluster Cluster -> Replica set Cluster -> Cluster

DDS Serving as the Source in Synchronization

Table 3-20 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
DDS	DDS Versions 3.2, 3.4, 4.0, 4.2, 4.4, and 5.0	<ul style="list-style-type: none"> Self-managed MongoDB Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 MongoDB on other clouds Versions 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, and 5.0 	Incremental synchronization supports the following modes: Replica set -> Replica set Cluster -> Cluster (the source cluster version must be 4.0 or later)
	DDS Versions 4.0, 4.2, 4.4, and 5.0	Kafka Version 0.11 or later	Incremental

DB2 for LUW Serving as the Source in Synchronization

Table 3-21 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
DB2 for LUW	DB2 for LUW Versions 9.7, 10.1, 10.5, 11.1, and 11.5	GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Full+Incremental
		GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Full+Incremental

TiDB Serving as the Source in Synchronization

Table 3-22 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
TiDB	TiDB Version 4.0.0 and later (excluding the development version)	TaurusDB Compatible with MySQL 8.0	Full+Incremental

Microsoft SQL Server as the Source in Synchronization

Table 3-23 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
Microsoft SQL Server	<ul style="list-style-type: none"> Self-managed Microsoft SQL Server Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 	GaussDB Distributed The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Incremental Full+Incremental
	Standard Edition 2016 SP2 or later, 2017, 2019 and 2022	GaussDB Centralized The database kernel version is 505.2 or earlier. Mapping Between Kernel Versions and DB Engine Versions.	Full Incremental Full+Incremental
	<ul style="list-style-type: none"> Microsoft SQL Server-compatible databases on other clouds Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 Standard Edition 2016 SP2 or later, 2017, 2019 and 2022 RDS for SQL Server Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 Standard Edition 2016 SP2 or later, 2017, 2019 and 2022 	<ul style="list-style-type: none"> Self-managed Microsoft SQL Server Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 Standard Edition 2016 SP2 or later, 2017, 2019 and 2022 Microsoft SQL Server-compatible databases on other clouds Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 Standard Edition 2016 SP2 or later, 2017, 2019 and 2022 RDS for SQL Server Enterprise Edition 2012, 2014, 2016, 2017, 2019 and 2022 Standard Edition 2016 SP2 or later, 2017, 2019 and 2022 	Full+Incremental Incremental

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
		Kafka Version 0.11 or later NOTE This function is in restricted use. To use it, submit a service ticket .	Incremental
		RDS for MySQL Versions 5.5, 5.6, 5.7, and 8.0 NOTE This function is in restricted use. To use it, submit a service ticket .	Full Incremental Full+Incremental
		TaurusDB Compatible with MySQL 8.0 NOTE This function is in restricted use. To use it, submit a service ticket .	Full Incremental Full+Incremental
		RDS for PostgreSQL Versions 9.5, 9.6, 10, 11, 12, 13, 14, 15, and 16 NOTE This function is in restricted use. To use it, submit a service ticket .	Full Incremental Full+Incremental

Cassandra Serving as the Source in Synchronization

Table 3-24 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
Cassandra	Self-managed Cassandra Version 2.0 or later	GeminiDB Cassandra Version 2.0 or later	Full

DynamoDB Serving as the Source in Synchronization

Table 3-25 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
DynamoDB	DynamoDB on other clouds (web services)	GeminiDB DynamoDB-Compatible API Version 3.11 NOTE This function is in restricted use. To use it, submit a service ticket .	Full Full+Incremental
GeminiDB DynamoDB-Compatible API	GeminiDB DynamoDB-Compatible API 3.11	DynamoDB on other clouds (web services) NOTE This function is in restricted use. To use it, submit a service ticket .	Full+Incremental Incremental

Redis Serving as the Source in Synchronization

Table 3-26 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
Redis	<ul style="list-style-type: none"> On-premises official open-source single-node or primary/standby Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) ECS-hosted official open-source single-node or primary/standby Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) 	GeminiDB Redis (classic deployment: Proxy Cluster and Primary/Standby; cloud native deployment: Proxy Cluster) NOTE The destination database version must be the same as or later than the source database version.	Full Full+Incremental

Redis Cluster Serving as the Source in Synchronization

Table 3-27 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
Redis Cluster	<ul style="list-style-type: none"> On-premises office open-source Redis Cluster (versions 3.0.x, 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) ECS-hosted official open-source Redis Cluster (versions 3.0.x, 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) 	GeminiDB Redis (classic deployment: Proxy Cluster and Primary/Standby; cloud native deployment: Proxy Cluster) NOTE The destination database version must be the same as or later than the source database version.	Full +Incremental

GeminiDB Redis Serving as the Source in Synchronization

Table 3-28 Database information

Source DB Engine	Source DB Type and Version	Destination DB Type and Version	Synchronization Mode
GeminiDB Redis	GeminiDB Redis (classic deployment : Proxy Cluster and Primary/Standby; cloud native deployment : Proxy Cluster) NOTE GeminiDB Redis performance-enhanced instances cannot be migrated out of the cloud.	<ul style="list-style-type: none"> On-premises Codis clusters (open-source Codis 3.0 or later) ECS-hosted Codis clusters (open-source Codis 3.0 or later) On-premises single-node Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) ECS-hosted single-node Redis (versions 2.8.x, 3.0.x, 3.2.x, 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) On-premises primary/standby Redis (versions 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) ECS-hosted primary/standby Redis (versions 4.0.x, 5.0.x, 6.0.x, 6.2.x, 7.0.x, and 7.2.x) GeminiDB Redis (classic deployment: Proxy Cluster Redis and Primary/Standby Redis) 	Incremental Full +Incremental
		<ul style="list-style-type: none"> ECS-hosted Redis Cluster (open-source Redis 4.0 or later) On-premises Redis Cluster (open-source Redis 4.0 or later) 	Incremental Full +Incremental

3.4 Data Subscription

Data subscription is used to obtain data changes of key services in the database. DRS caches the data changes and provides a unified SDK interface for downstream services to subscribe to, obtain, and consume the changes. [Table 3-29](#) describes the supported databases and data types. For details about data subscription, see [Data Subscription](#).

Table 3-29 Database information

DB Engine	Data Type
RDS for MySQL Versions 5.6 and 5.7	<ul style="list-style-type: none"> Data update Structure update

3.5 Real-Time Disaster Recovery

To prevent service unavailability caused by regional faults, DRS provides disaster recovery to ensure service continuity. The following table lists the databases and versions supported by real-time DR. For more information about real-time DR, see [Real-Time Disaster Recovery](#).

Self-managed databases (for example, MySQL) include on-premises databases and the databases created on ECSs. RDS for MySQL refers to the MySQL databases created using Huawei Cloud RDS.

NOTE

- Some functions are in restricted use. To use them, [submit a service ticket](#).
- The destination database version must be the same as or later than the source database version.

MySQL Serving as the Source in DR

Table 3-30 Database information

Service DB Engine	Service DB Type and Version	DR DB Type and Version
MySQL	<ul style="list-style-type: none"> • Self-managed MySQL Versions 5.6, 5.7, and 8.0 • MySQL on other clouds Versions 5.6, 5.7, and 8.0 	RDS for MySQL All versions
		TaurusDB Compatible with MySQL 8.0
	RDS for MySQL All versions	RDS for MySQL All versions
		<ul style="list-style-type: none"> • Self-managed MySQL Versions 5.6, 5.7, and 8.0 • MySQL on other clouds Versions 5.6, 5.7, and 8.0
		TaurusDB Compatible with MySQL 8.0

DDM Serving as the Source in DR

Table 3-31 Database information

Service DB Engine	Service DB Type and Version	DR DB Type and Version
DDM	DDM	DDM

TaurusDB Serving as the Source in DR

Service DB Engine	Service DB Type and Version	DR DB Type and Version
TaurusDB	TaurusDB Compatible with MySQL 8.0	TaurusDB Compatible with MySQL 8.0

3.6 Workload Replay

A workload replay task simulates the source database's workload and runs the workload on the destination database, so you can evaluate the functionality and performance of the destination database. The following table lists the supported databases and versions. For more information about workload replay, see [Workload Replay](#).

Self-managed databases (such as MySQL and MariaDB) include on-premises databases and databases created on ECSs. RDS for MySQL refers to the MySQL databases created using Huawei Cloud RDS.

NOTE

- The destination database version must be the same as or later than the source database version.

Table 3-32 Database information

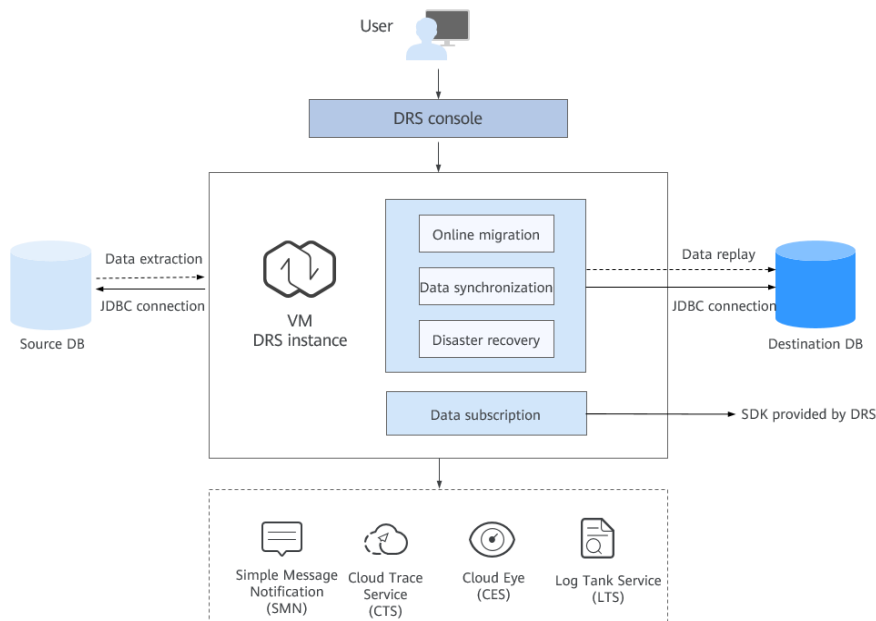
Source DB Engine	Source DB Type and Version	Destination DB Type and Version
MySQL	RDS for MySQL All versions	RDS for MySQL All versions
		TaurusDB Compatible with MySQL 8.0
	<ul style="list-style-type: none"> Self-managed MySQL Versions 5.6, 5.7, and 8.0 MySQL on other clouds Versions 5.6, 5.7, and 8.0 	RDS for MySQL All versions
		TaurusDB Compatible with MySQL 8.0
TaurusDB	TaurusDB Compatible with MySQL 8.0	TaurusDB Compatible with MySQL 8.0

Source DB Engine	Source DB Type and Version	Destination DB Type and Version
MongoDB	<ul style="list-style-type: none"> On-premises MongoDB 5.0 On-premises MongoDB 6.0 On-premises MongoDB 7.0 	<ul style="list-style-type: none"> DDS 4.4 DDS 5.0
MariaDB	<ul style="list-style-type: none"> Self-managed MariaDB Versions 10.3, 10.4, and 10.5 MariaDB on other clouds Versions 10.3, 10.4, and 10.5 	RDS for MariaDB Version 10.5
MariaDB	<ul style="list-style-type: none"> Self-managed MariaDB Versions 10.3, 10.4, and 10.5 MariaDB on other clouds Versions 10.3, 10.4, and 10.5 	RDS for MySQL All versions <ul style="list-style-type: none"> TaurusDB Basic Edition 5.7 and 8.0
MariaDB	<ul style="list-style-type: none"> Self-managed MariaDB Versions 10.3, 10.4, and 10.5 MariaDB on other clouds Versions 10.3, 10.4, and 10.5 	TaurusDB Compatible with MySQL 8.0

4 Product Architecture and Function Principles

The following figure shows the product architecture and function principles of DRS.

Figure 4-1 DRS product architecture



Architecture Description

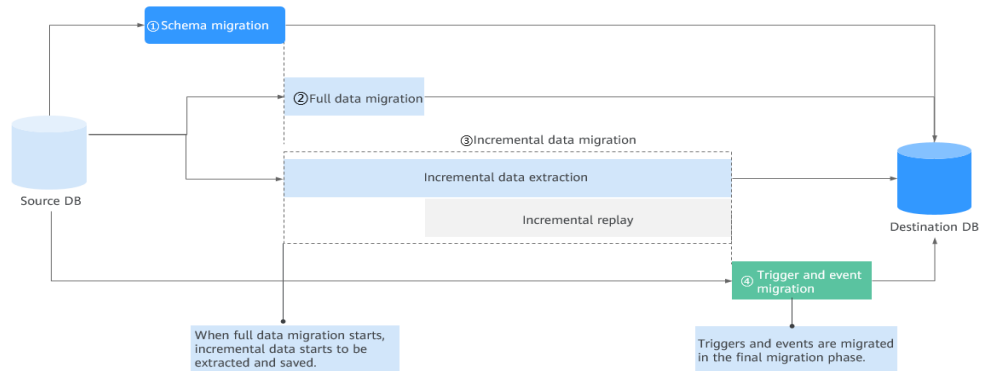
- Minimum permission design
 - a. Java Database Connectivity (JDBC) is used to connect to the source and destination databases, so you do not have to deploy programs on the databases.
 - b. A task runs on an independent and exclusively used VM. Data is isolated between tenants.
 - c. The number of IP addresses is limited. Only the DRS instance IP address is allowed to access the source and destination databases.

- Reliability design
 - a. Automatic reconnection: If the connection between DRS and your database breaks down due to bad network or database switchover, DRS automatically retries the connection until the task is restored.
 - b. Resumable upload: When the connection to the source or the destination is abnormal, DRS automatically marks the current replay point. After the fault is rectified, you can resume data transfer from the replay point to ensure data consistency.
 - c. If the VM where the DRS replication instance is located fails, services are automatically switched to a new VM with the IP address unchanged to ensure that the migration task is not interrupted.

The character set standard used by DRS is Unicode 6.2.0.

Basic Principles of Real-Time Migration

Figure 4-2 Real-time migration principle




- Take the full+incremental migration as an example. A complete migration process includes four phases.
 - a. **Phase 1:** Structure migration. DRS queries the databases, tables, and primary keys to be migrated from the source and creates corresponding objects in the destination.
 - b. **Phase 2:** Full data migration. DRS uses the parallel technology to query all data from the source and inserts the data into the destination, which is fast and convenient. Before the full migration is started, incremental data is extracted and saved in advance to ensure data integrity and consistency in the subsequent incremental migration process.
 - c. **Phase 3:** Incremental data migration. After the full migration task is complete, the incremental migration task is started. The incremental data generated after the start of the full migration is continuously parsed, converted, and replayed to the destination database until data is in sync between the source and destination databases.
 - d. **Phase 4:** To prevent data from being operated by triggers and events during the migration, triggers and events will be migrated after a migration task is complete.

- Principles of the underlying module for full migration:
 - Sharding module:** calculates the sharding logic of each table using the optimized sharding algorithm.
 - Extraction module:** queries data from the source database in parallel mode based on the calculated shard information.
 - Replay module:** inserts the data queried by the extraction module into the destination database in parallel and multi-task mode.
- Principles of the underlying module for incremental migration:
 - Log reading module:** reads the original incremental log data (for example, binlog for MySQL) from the source database, parses the data, converts the data into the standard log format, and stores it locally.
 - Log replay module:** processes and filters incremental logs based on the standard format converted by the log reading module, and synchronizes the incremental data to the destination database.

Basic Principles of Backup Migration

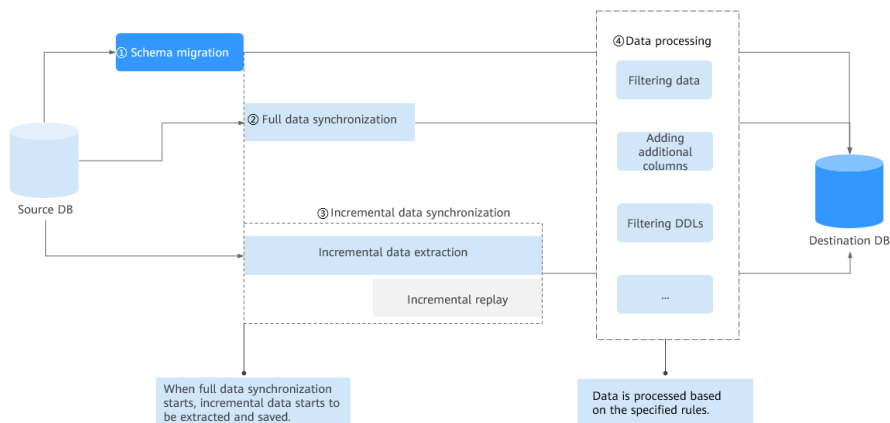
Figure 4-3 Backup migration principle



DRS allows you to migrate data from a Microsoft SQL Server database to the cloud using the backup file of the database. You can copy the full and incremental backup files of the source database to an OBS bucket. DRS downloads that files from the bucket and uploads them to the disk of the destination database. After the pre-check and verification are complete, DRS runs the import command to restore the data to the destination database.

Basic Principles of Real-Time Synchronization

Figure 4-4 Real-time synchronization principle

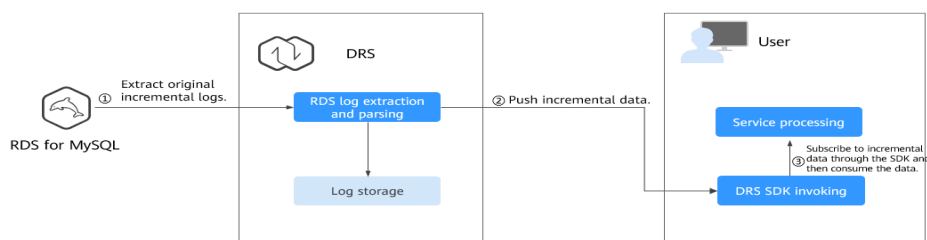


Real-time synchronization can ensure that data is always in sync between the source and destination databases. It mainly applies to synchronization from OLTP to OLAP or from OLTP to big data components in real time. The technical principles of full+incremental synchronization and real-time migration are basically the same. However, there is a slight difference between them in different scenarios.

1. DRS supports heterogeneous synchronization (between different DB engines). It means that DRS converts the structure definition statements of the source database to match that of the destination database. In addition, DRS can map and convert database field types. You can refer to [Mapping Data Types](#) of heterogeneous databases or use Database and Application Migration UGO (UGO) to synchronize the structure of heterogeneous databases.
2. DRS allows you to configure data processing rules, so you can use these rules to extract, parse, and replay data to meet your service requirements.
3. Objects such as accounts, triggers, and events cannot be synchronized.
4. Real-time synchronization is often used in many-to-one scenario. DDL operations in many-to-one and one-to-many scenarios are specially processed.

Basic Principles of Data Subscription

Figure 4-5 Data subscription principle



Data subscription provides an SDK so that customers' service programs can obtain incremental data from the source database in real time.

DRS extracts original incremental logs from the source database, parses the logs into the standard format, and persists the logs to the local host. In addition, DRS invokes the notification interface of the client subscription SDK in real time to push incremental data to the client service program. Then, the client can consume data changes based on service requirements.

The incremental data consumed by the client program is recorded on the server in real time. The DRS server can continue to push incremental data from the last consumption position in scenarios such as service interruption and reconnection.

Basic Principles of Real-Time Disaster Recovery

DRS uses the real-time replication technology to implement disaster recovery for two databases. The underlying technical principles are the same as those of real-time migration. The difference is that real-time DR supports forward

synchronization and backward synchronization. In addition, disaster recovery is performed on the instance-level, which means that databases and tables cannot be selected.

Basic Principles of Workload Replay

A workload replay task consists of SQL recording and replay. All of the SQL statements (create, delete, update, and query operations) executed in the required period on the source database will be downloaded by a recording tool from the binlog, and then cached and injected into the destination database where you can trigger a replay and review performance. By specifying the replay thread and speed, you can simulate the peak service load of the source database and analyze the stability of the destination database when workloads increase sharply.

5 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, and supports on-demand pricing.

Secure

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

6 Functions and Features

6.1 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.

 NOTE

Some functions are in restricted use. To use them, [submit a service ticket](#).

Supported Database Types

For details about the supported databases, versions, and migration types, see [Supported Databases](#).

Supported Network Types

DRS supports data migration through a Virtual Private Cloud (VPC), Virtual Private Network (VPN), Direct Connect, or public network. [Table 6-1](#) lists the application scenarios of each network type and required preparations, and [Table 6-2](#) lists the supported network types of each migration scenario.

Table 6-1 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 does not support communication between the source database and destination database over a VPC across tenants. If necessary, you can create a VPC peering connection and select VPN for Network Type to enable communication between the source and destination databases. For details about how to create a VPC peering connection, see VPC Peering Connection in the <i>Virtual Private Cloud User Guide</i>.
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. For more information about VPN, see Getting Started with Virtual Private Network .
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. For more information about Direct Connect, see Getting Started with Direct Connect .

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 6-2 Supported network types

Migration Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
To the cloud	MySQL	MySQL	Supported	Supported	Supported
		DDM	Supported	Supported	Supported
		TaurusDB	Supported	Supported	Supported
	MongoDB	DDS	Supported	Supported	Supported
		GeminiDB Mongo	Supported	Supported	Supported

Migration Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
	MySQL schema and logic table	DDM	Supported	Supported	Supported
	Redis	GeminiDB Redis	Supported	Supported	Supported
	Redis Cluster	GeminiDB Redis	Supported	Supported	Supported
From the cloud	MySQL	MySQL	Supported	Supported	Supported
	DDS	MongoDB	Supported	Supported	Supported
	GeminiDB Redis	Redis	Supported	Supported	Supported
		Redis Cluster	Supported	Supported	Supported

Migration Objects

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

Table 6-3 Supported migration objects

Direction	Source DB	Destination DB	Full Migration	Table-Level Migration	Database-Level Migration
To the cloud	MySQL	MySQL	Supported	Supported	Supported
		TaurusDB	Supported	Supported	Supported
		DDM	Not supported	Supported	Not supported
	MongoDB	DDS	Supported	Supported	Supported
		GeminiDB Mongo	Supported	Supported	Supported

Direction	Source DB	Destination DB	Full Migration	Table-Level Migration	Database-Level Migration
	MySQL schema and logic table	DDM	Not supported	Supported	Not supported
	Redis	GeminiDB Redis	Not supported	Not supported	Supported
	Redis Cluster	GeminiDB Redis	Supported	Not supported	Not supported
From the cloud	MySQL	MySQL	Supported	Supported	Supported
	DDS	MongoDB	Supported	Supported	Supported
	GeminiDB Redis	Redis	Supported	Not supported	Not supported
	GeminiDB Redis	Redis Cluster	Supported	Not supported	Not supported

Advanced Features

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

Table 6-4 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.

6.2 Backup Migration

DRS supports backup migrations of various database types.

Supported Database Types

Table 6-5 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 files RDS for SQL Server full backup files Microsoft SQL Server backup files on other clouds 	RDS for SQL Server DB instances

Migration Methods

Table 6-6 Migration methods

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

Supported Database Versions

Table 6-7 Database versions

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

Data Flow	Backup File Version	Destination DB Version
	<p>RDS for SQL Server full backup file versions:</p> <ul style="list-style-type: none"> • Microsoft SQL Server 2008 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2012 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2014 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2016 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2017 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2019 Enterprise Edition, Standard Edition, and Web Edition 	<ul style="list-style-type: none"> • Microsoft SQL Server 2008 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2012 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2014 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2016 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2017 Enterprise Edition, Standard Edition, and Web Edition • Microsoft SQL Server 2019 Enterprise Edition, Standard Edition, and Web Edition

Data Flow	Backup File Version	Destination DB Version
	<p>RDS for SQL Server full backup file versions:</p> <ul style="list-style-type: none">• Microsoft SQL Server 2012 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2014 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2016 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2017 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2019 Enterprise Edition, Standard Edition, and Web Edition• Microsoft SQL Server 2022 Enterprise Edition, Standard Edition, and Web Edition	<ul style="list-style-type: none">• Microsoft SQL Server 2022 Enterprise Edition, Standard Edition, and Web Edition

Backup Migration Scenarios

Table 6-8 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.

6.3 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.

NOTE

Some functions are in restricted use. To use them, [submit a service ticket](#).

Supported Database Types

For details about supported database types and versions and synchronization types, see [Supported Databases](#).

Network Types

DRS supports real-time synchronization through a Virtual Private Cloud (VPC), Virtual Private Network (VPN), Direct Connect, or public network. [Table 6-9](#) lists the application scenarios of each network type and required preparations, and [Table 6-10](#) lists the supported network types of each synchronization scenario.

Table 6-9 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 does not support communication between the source database and destination database over a VPC across tenants. If necessary, you can create a VPC peering connection and select VPN for Network Type to enable communication between the source and destination databases. For details about how to create a VPC peering connection, see VPC Peering Connection in the <i>Virtual Private Cloud User Guide</i>.
VPN	Synchronization from on-premises databases to cloud databases or between cloud databases across regions	<p>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.</p> <p>For more information about VPN, see Getting Started with Virtual Private Network.</p>

Network Type	Application Scenario	Preparations
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. For more information about Direct Connect, see Getting Started with Direct Connect .
Public network	Synchronization from on-premises or external cloud databases to the destination databases.	To ensure network connectivity between the source and destination databases, perform the following operations: <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 6-10 Supported network types

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
To the cloud	MySQL	MySQL	Supported	Supported	Supported

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
		PostgreSQL	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
		TaurusDB	Supported	Supported	Supported
		MariaDB	Supported	Supported	Supported
	PostgreSQL	PostgreSQL	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
	DDM	MySQL	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
		DDM	Supported	Supported	Supported
	Oracle	MySQL	Supported	Supported	Supported
		DDM	Supported	Supported	Supported
		PostgreSQL	Supported	Supported	Supported
		TaurusDB	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
		GaussDB Distributed	Supported	Supported	Supported
	DB2 for LUW	GaussDB Centralized	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
	TiDB	TaurusDB	Not supported	Supported	Supported
	Microsoft SQL Server	GaussDB Centralized	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
		Microsoft SQL Server	Supported	Supported	Supported
		MySQL NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Supported	Supported
		TaurusDB NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Supported	Supported
		PostgreSQL NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Supported	Supported
	MongoDB	DDS	Supported	Supported	Supported
	TaurusDB	TaurusDB	Supported	Supported	Supported
	MariaDB	MariaDB	Supported	Supported	Supported

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
		MySQL	Supported	Supported	Supported
		TaurusDB	Supported	Supported	Supported
	DynamoDB	GeminiDB DynamoDB- Compatible API NOTE This function is in restricted use. To use it, submit a service ticket .	Not supported	Supported	Supported
	Redis	GeminiDB Redis	Supported	Supported	Supported
	Redis Cluster	GeminiDB Redis	Supported	Supported	Supported
	Cassandra	GeminiDB Cassandra	Supported	Supported	Supported
From the cloud	MySQL	MySQL	Supported	Supported	Supported
		Kafka	Supported	Supported	Supported
		Oracle	Supported	Supported	Supported
		MariaDB	Supported	Supported	Supported
	DDM	MySQL	Supported	Supported	Supported
		Oracle	Supported	Supported	Supported
		Kafka	Supported	Supported	Supported
	DDS	MongoDB	Supported	Supported	Supported
		Kafka	Supported	Supported	Supported

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
	PostgreSQL	PostgreSQL	Supported	Supported	Supported
		Kafka	Not supported	Supported	Supported
	GaussDB Centralized	MySQL	Not supported	Supported	Supported
		Oracle	Not supported	Supported	Supported
		Kafka	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
	GaussDB Distributed	MySQL	Not supported	Supported	Supported
		Oracle	Not supported	Supported	Supported
		Kafka	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
	TaurusDB	MySQL	Supported	Supported	Supported
		Kafka	Supported	Supported	Supported
		Oracle	Supported	Supported	Supported

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
	MariaDB	MariaDB	Supported	Supported	Supported
	GeminiDB Redis	Redis	Supported	Supported	Supported
		Redis Cluster	Supported	Supported	Supported
	Microsoft SQL Server	Kafka NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Supported	Supported
Self-built -> Self-built	MySQL	Kafka	Supported	Supported	Supported
		GaussDB Centralized	Not supported	Supported	Supported
		GaussDB Distributed	Not supported	Supported	Supported
	Oracle	Kafka	Supported	Supported	Supported
		GaussDB Centralized	Not supported	Supported	Supported
		GaussDB Distributed	Not supported	Supported	Supported
	PostgreSQL	Kafka	Not supported	Supported	Supported
	GaussDB Centralized	MySQL	Not supported	Supported	Supported
		Oracle	Not supported	Supported	Supported

Synchronization Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
		Kafka	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
	GaussDB Distributed	MySQL	Not supported	Supported	Supported
		Oracle	Not supported	Supported	Supported
		Kafka	Supported	Supported	Supported
		GaussDB Distributed	Supported	Supported	Supported
		GaussDB Centralized	Supported	Supported	Supported
	DB2 for LUW	GaussDB Centralized	Not supported	Supported	Supported
		GaussDB Distributed	Not supported	Supported	Supported
	Microsoft SQL Server	Kafka NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Supported	Supported

Supported Synchronization Objects

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

Table 6-11 Supported synchronization objects

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File
To the cloud	MySQL	MySQL	Supported	Supported	Supported
		PostgreSQL	Supported	Supported	Supported
		GaussDB Distributed	Supported	Not supported	Supported
		GaussDB Centralized	Supported	Not supported	Supported
		TaurusDB	Supported	Supported	Supported
		MariaDB	Supported	Supported	Supported
	PostgreSQL	PostgreSQL	Supported	Supported	Supported
		GaussDB Centralized	Supported	Not supported	Supported
		GaussDB Distributed	Supported	Not supported	Supported
	DDM	MySQL	Supported	Not supported	Not supported
		GaussDB Distributed	Supported	Not supported	Supported
		GaussDB Centralized	Supported	Not supported	Supported
		DDM	Supported	Supported	Not supported
	Oracle	MySQL	Supported	Not supported	Supported

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File
		DDM	Supported	Not supported	Not supported
		PostgreSQL	Supported	Not supported	Supported
		TaurusDB	Supported	Not supported	Supported
		GaussDB Centralized	Supported	Not supported	Supported
		GaussDB Distributed	Supported	Not supported	Supported
	DB2 for LUW	GaussDB Centralized	Supported	Not supported	Supported
		GaussDB Distributed	Supported	Not supported	Supported
	TiDB	TaurusDB	Supported	Supported	Not supported
	Microsoft SQL Server	GaussDB Centralized	Supported	Not supported	Not supported
		GaussDB Distributed	Supported	Not supported	Not supported
		Microsoft SQL Server	Supported	Not supported	Not supported
		MySQL NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Not supported	Supported

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File
		TaurusDB NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Not supported	Supported
		PostgreSQL NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Not supported	Supported
	MongoDB	DDS	Supported	Supported	Not supported
	TaurusDB	TaurusDB	Supported	Supported	Supported
	MariaDB	MariaDB	Supported	Supported	Supported
		MySQL	Supported	Supported	Supported
		TaurusDB	Supported	Supported	Supported
	DynamoDB	GeminiDB DynamoDB- Compatible API NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Not supported	Not supported
	Cassandra	GeminiDB Cassandra	Supported	Not supported	Supported
	From the cloud	MySQL	MySQL	Supported	Supported
Kafka			Supported	Supported	Supported
Oracle			Supported	Not supported	Supported

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File
		MariaDB	Supported	Supported	Supported
	DDM	MySQL	Supported	Not supported	Not supported
		Oracle	Supported	Supported	Not supported
		Kafka	Supported	Supported	Not supported
	DDS	MongoDB	Supported	Supported	Not supported
		Kafka	Supported	Supported	Not supported
	PostgreSQL	PostgreSQL	Supported	Supported	Supported
		Kafka	Supported	Supported	Not supported
	GaussDB Centralized	MySQL	Supported	Not supported	Not supported
		Oracle	Supported	Not supported	Supported
		Kafka	Supported	Supported	Not supported
		GaussDB Distributed	Supported	Not supported	Supported
		GaussDB Centralized	Supported	Not supported	Supported
	GaussDB Distributed	MySQL	Supported	Not supported	Not supported
		Oracle	Supported	Not supported	Supported

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File	
		Kafka	Supported	Supported	Not supported	
		GaussDB Distributed	Supported	Not supported	Supported	
		GaussDB Centralized	Supported	Not supported	Supported	
	TaurusDB	MySQL	Supported	Supported	Supported	
		Kafka	Supported	Supported	Supported	
		Oracle	Supported	Not supported	Supported	
	MariaDB	MariaDB	Supported	Supported	Supported	
	Microsoft SQL Server	Kafka NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Not supported	Supported	
	Self-built -> Self-built	MySQL	Kafka	Supported	Supported	Supported
			GaussDB Centralized	Supported	Not supported	Supported
GaussDB Distributed			Supported	Not supported	Supported	
Oracle		Kafka	Supported	Not supported	Supported	
		GaussDB Centralized	Supported	Not supported	Supported	

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File
		GaussDB Distributed	Supported	Not supported	Supported
	GaussDB Centralized	MySQL	Supported	Not supported	Not supported
		Oracle	Supported	Not supported	Supported
		Kafka	Supported	Supported	Not supported
		GaussDB Centralized	Supported	Not supported	Supported
		GaussDB Distributed	Supported	Not supported	Supported
	GaussDB Distributed	MySQL	Supported	Not supported	Not supported
		Oracle	Supported	Not supported	Supported
		Kafka	Supported	Supported	Not supported
		GaussDB Distributed	Supported	Not supported	Supported
		GaussDB Centralized	Supported	Not supported	Supported
	PostgreSQL	Kafka	Supported	Supported	Not supported
	DB2 for LUW	GaussDB Centralized	Supported	Not supported	Supported

Synchronization Direction	Source DB	Destination DB	Table-level	Database-level	Importing an Object File
		GaussDB Distributed	Supported	Not supported	Supported
	Microsoft SQL Server	Kafka NOTE This function is in restricted use. To use it, submit a service ticket .	Supported	Not supported	Supported

Advanced Features

DRS supports multiple features to ensure successful data synchronization.

Table 6-12 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	<p>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.</p>

6.4 Data Subscription

Data Source Types

Data subscription supports the following data source type:

- RDS for MySQL DB instances

Subscription Objects

The subscription objects are tables.

The incremental data is divided into Data Manipulation Language (DML) and Data Definition Language (DDL).

Database Versions

DRS supports the following source database versions:

- MySQL 5.6.x
- MySQL 5.7.x

Advanced Features

Data subscription provides multiple features. For details, see [Table 6-13](#).

Table 6-13 Advanced features

Feature	Description
Dynamically adding or deleting subscription objects	During data subscription, you add or delete subscription objects as required.
Viewing subscription data	View the incremental data on the management console.
Modifying the consumption start time	During consumption, you can change the consumption start time at any time.

6.5 Real-Time Disaster Recovery

The following table lists the supported databases and versions.

 **NOTE**

Some functions are in restricted use. To use them, [submit a service ticket](#).

Network Preparations

DRS supports disaster recovery through a Virtual Private Network (VPN), Direct Connect, or public network. [Table 6-14](#) lists the application scenarios of each network type and required preparations, and [Table 6-15](#) lists the supported network types of each DR scenario.

Table 6-14 Network types

Network Type	Application Scenario	Preparations
VPN	Disaster recovery 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 disaster recovery, ensure that the VPN network is accessible. For more information about VPN, see Getting Started with Virtual Private Network .

Network Type	Application Scenario	Preparations
Direct Connect	Disaster recovery 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. For more information about Direct Connect, see Getting Started with Direct Connect .
Public network	Disaster recovery from on-premises databases or other cloud databases to destination databases.	To ensure network connectivity between the source and destination databases, perform the following operations: <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 disaster recovery instance to the whitelist of the source database for inbound traffic. • If destination databases and the DR 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 DR instance. • If SSL is not enabled, backing up confidential data for disaster recovery is not recommended.

Table 6-15 Supported network types

DR Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
Current cloud as standby	MySQL	MySQL	Not supported	Supported	Supported
		TaurusDB	Not supported	Supported	Supported

DR Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
	DDM	DDM	Not supported	Supported	Supported
	TaurusDB	TaurusDB	Not supported	Supported	Supported
Current cloud as active	MySQL	MySQL	Not supported	Supported	Supported
	DDM	DDM	Not supported	Supported	Supported
	TaurusDB	TaurusDB	Not supported	Supported	Supported
Dual-Active DR	MySQL NOTE This function is in restricted use.	MySQL	Not supported	Supported	Supported
	TaurusDB NOTE This function is in restricted use.	TaurusDB	Not supported	Supported	Supported

6.6 Workload Replay

The following table lists the supported databases and versions.

Database Types

DRS supports workload replay management for the following types of databases.

Table 6-16 Database types

Replay Direction	Data Flow	Source DB	Destination DB	Destination DB Type
Current cloud	MySQL->MySQL	RDS for MySQL instances	RDS for MySQL instances	<ul style="list-style-type: none"> • Single • Primary / Standby
Current cloud	MySQL->TaurusDB	RDS for MySQL instances	TaurusDB instances	Primary/ Standby
Current cloud	TaurusDB->TaurusDB	TaurusDB instances	TaurusDB instances	Primary/ Standby
Current cloud	TaurusDB->MySQL	TaurusDB instances	RDS for MySQL instances	<ul style="list-style-type: none"> • Single • Primary / Standby
To the cloud	MySQL->MySQL	<ul style="list-style-type: none"> • On-premises MySQL databases • ECS-hosted MySQL databases • MySQL databases on other clouds 	RDS for MySQL instances	<ul style="list-style-type: none"> • Single • Primary / Standby

Replay Direction	Data Flow	Source DB	Destination DB	Destination DB Type
To the cloud	MySQL->TaurusDB	<ul style="list-style-type: none"> On-premises MySQL databases ECS-hosted MySQL databases MySQL databases on other clouds 	TaurusDB instances	Primary/Standby
To the cloud	Self-built MongoDB -> DDS	On-premises MongoDB	DDS instances	<ul style="list-style-type: none"> Cluster Replica set Single node
To the cloud	MariaDB->MySQL	<ul style="list-style-type: none"> On-premises MariaDB ECS-hosted MariaDB MariaDB on other clouds 	RDS for MySQL instances <ul style="list-style-type: none"> Basic edition TaurusDB instances 	<ul style="list-style-type: none"> Single Primary / Standby
To the cloud	MariaDB->MariaDB	<ul style="list-style-type: none"> On-premises MariaDB ECS-hosted MariaDB MariaDB on other clouds 	RDS for MariaDB instances	<ul style="list-style-type: none"> Single Primary / Standby

Replay Direction	Data Flow	Source DB	Destination DB	Destination DB Type
To the cloud	MariaDB->TaurusDB	<ul style="list-style-type: none"> On-premises MariaDB ECS-hosted MariaDB MariaDB on other clouds 	TaurusDB instances	Primary/Standby

Database Versions

Table 6-17 Database versions

Replay Direction	Data Flow	Source DB Version	Destination DB Version
Current 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
Current cloud	MySQL->TaurusDB	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x 	Compatible with MySQL 8.0
Current cloud	TaurusDB->TaurusDB	Compatible with MySQL 8.0	Compatible with MySQL 8.0
Current cloud	TaurusDB->MySQL	Compatible with MySQL 8.0	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x

Replay Direction	Data Flow	Source DB Version	Destination DB Version
To 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
To the cloud	MySQL->TaurusDB	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x 	Compatible with MySQL 8.0
To the cloud	Self-built MongoDB -> DDS	<ul style="list-style-type: none"> On-premises MongoDB 5.0 On-premises MongoDB 6.0 On-premises MongoDB 7.0 	<ul style="list-style-type: none"> DDS 4.4 DDS 5.0
To the cloud	MariaDB->MySQL	<ul style="list-style-type: none"> MariaDB 10.3 MariaDB 10.4 MariaDB 10.5 	<ul style="list-style-type: none"> MySQL 5.6.x MySQL 5.7.x MySQL 8.0.x
To the cloud	MariaDB->MariaDB	<ul style="list-style-type: none"> MariaDB 10.3 MariaDB 10.4 MariaDB 10.5 	<ul style="list-style-type: none"> MariaDB 10.5
To the cloud	MariaDB->TaurusDB	<ul style="list-style-type: none"> MariaDB 10.3 MariaDB 10.4 MariaDB 10.5 	Compatible with MySQL 8.0

Network Types

DRS supports workload replay through a Virtual Private Cloud (VPC), Virtual Private Network (VPN), Direct Connect, or public network. [Table 6-18](#) lists the application scenarios of each network type and required preparations, and [Table 6-19](#) lists the supported network types of each workload replay scenario.

Table 6-18 Network types

Network Type	Application Scenario	Preparations
VPC	Workload replay 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. 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 does not support communication between the source database and destination database over a VPC across tenants. If necessary, you can create a VPC peering connection and select VPN for Network Type to enable communication between the source and destination databases. For details about how to create a VPC peering connection, see Virtual Private Cloud User Guide.
VPN	Workload replay from on-premises databases to cloud databases or between cloud databases across regions	<p>Establish a VPN connection between your local data center and the VPC that hosts the destination database. Before workload replay, ensure that the VPN network is accessible.</p> <p>For more information about VPN, see Getting Started with Virtual Private Network.</p>
Direct Connect	Workload replay from on-premises databases to cloud databases or between cloud databases across regions	<p>Use a dedicated network connection to connect your data center to VPCs.</p> <p>For more information about Direct Connect, see Getting Started with Direct Connect.</p>

Network Type	Application Scenario	Preparations
Public network	Workload replay from on-premises databases 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 DRS instance to the whitelist of the source database for inbound traffic. • If destination databases and the DRS 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 displayed on the Configure Source and Destination Databases page is the EIP of the DRS instance.

Table 6-19 Supported network types

Replay Direction	Source DB	Destination DB	VPC	Public Network	VPN or Direct Connect
Current cloud	MySQL	MySQL	Supported	Supported	Supported
Current cloud	MySQL	TaurusDB	Supported	Supported	Supported
Current cloud	TaurusDB	TaurusDB	Supported	Supported	Supported

Rep lay Dir ecti on	Source DB	Destination DB	VPC	Publi c Netw ork	VPN or Direct Connect
Cur rent clo ud	TaurusDB	MySQL	Suppo rted	Suppo rted	Supporte d
To the clo ud	MySQL	MySQL	Suppo rted	Suppo rted	Supporte d
To the clo ud	MySQL	TaurusDB	Suppo rted	Suppo rted	Supporte d
To the clo ud	Self-built MongoDB	DDS	Suppo rted	Suppo rted	Supporte d
To the clo ud	MariaDB	MySQL	Suppo rted	Suppo rted	Supporte d
To the clo ud	MariaDB	MariaDB	Suppo rted	Suppo rted	Supporte d
To the clo ud	MariaDB	TaurusDB	Suppo rted	Suppo rted	Supporte d

7 Specification Description

7.1 Real-Time Synchronization

Precautions

The performance indicators provided in this section are for reference only. The actual environment is affected by factors such as the performance of the source or destination database, network bandwidth, data model, and service model.

Specifications

Table 7-1 Full synchronization maximum performance

Data Flow	Reference Value of Maximum Performance of Full Synchronization (MB/s)
MySQL as the source	50
Oracle as the source	40
Redis as the source	30
GaussDB as the source	40
PostgreSQL as the source	30
DDM as the source	20
MongoDB as the source	20

 **CAUTION**

- There are many factors that affect the DRS migration speed. The current migration speed is the test data when **there is no network and database performance bottlenecks and the task specifications are large**. The migration speed is for reference only.
- When the destination database is Oracle, the migration speed in the full phase is 30% to 50% lower than that of other types of databases due to the write mechanism of the destination database.
- The write performance of the MongoDB database is affected by the number of indexes. The write performance decreases by 5% to 8% when there is one index. The more indexes, the slower the write speed.

Based on the maximum incremental performance of data flows, there are six types of specifications: micro, small, medium, large, ultra-large, and macro. [Table 7-2](#) lists the maximum performance of each specification.

Table 7-2 Maximum performance

Specifications	Reference Value of Maximum Performance of Incremental Synchronization (Rows/Second)
Micro	300
Small	3,000
Medium	7,500
Large	10,000
Ultra-large	20,000
Macro	> 20,000

 **NOTE**

- The performance of each specification is affected by factors such as the networks, source and destination database performance, and latency. The values in the table are for reference only.
- DRS measures the performance of different specifications using the full (with flow control disabled) and incremental synchronization tasks as the standard.
- The maximum performance (rows/second) is measured by the number of rows synchronized per second, including every row of data changes caused by DML statements (INSERT, DELETE, and UPDATE). You can [view the destination database write frequency \(apply_rows_rate\) monitoring metric](#) on Cloud Eye.
- DRS allows you to upgrade specifications only for single-AZ synchronization tasks. Task specifications cannot be upgraded for dual-AZ tasks or downgraded. For details, see [Changing Specifications](#).
- If you want to compare values for a DRS task, select large or higher specifications when creating the DRS task.

Testing Models

Create a full+incremental real-time synchronization task for two RDS for MySQL instances. [Table 7-3](#) shows the instance configurations.

Table 7-3 Instance specifications

Parameter	Source RDS for MySQL instance	Destination RDS for MySQL instance
Flavor	c6.4xlarge.4 (general-enhanced II)	c6.4xlarge.4 (general-enhanced II)
Instance specifications	Ultra-high I/O	Ultra-high I/O
Storage type	16 vCPUs 64 GB	16 vCPUs 64 GB
Storage space	300 GB	300 GB
Maximum connections	18,000	18,000
Maximum QPS	3,325	3,325
Maximum IOPS	114,152	114,152

Test model:

- The number of test tables is 20.
- All test tables have primary keys.
- The record size is 1 KB.
- Each transaction contains two DML operations and one COMMIT operation. The ratio of INSERT, UPDATE, and DELETE operations is 1:1:1.

Multiple Specifications

DRS real-time synchronization allows you to select specifications for some specified data flow tasks.

Table 7-4 Data flow types that support multiple specifications

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
To the cloud	MySQL -> MySQL	Supported	Supported only for single-AZ tasks
To the cloud	MySQL -> PostgreSQL	Supported	Supported only for single-AZ tasks

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
To the cloud	MySQL -> GaussDB Distributed	Supported	Supported only for single-AZ tasks
To the cloud	MySQL -> GaussDB Centralized	Supported	Supported only for single-AZ tasks
To the cloud	MySQL -> TaurusDB	Supported	Supported only for single-AZ tasks
To the cloud	MySQL -> MariaDB	Not supported	Not supported
To the cloud	PostgreSQL -> PostgreSQL	Supported	Supported only for single-AZ tasks
To the cloud	PostgreSQL -> GaussDB Centralized	Supported	Supported
To the cloud	PostgreSQL -> GaussDB Distributed	Not supported	Not supported
To the cloud	DDM -> MySQL	Supported	Not supported
To the cloud	DDM -> DDM	Supported	Not supported
To the cloud	Oracle -> PostgreSQL	Supported	Supported only for single-AZ tasks
To the cloud	Oracle -> MySQL	Supported	Supported only for single-AZ tasks
To the cloud	Oracle -> TaurusDB	Supported	Supported only for single-AZ tasks

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
To the cloud	Oracle -> GaussDB Centralized	Supported	Supported only for single-AZ tasks
To the cloud	Oracle -> GaussDB Distributed	Supported	Supported only for single-AZ tasks
To the cloud	Oracle -> DDM	Supported	Supported only for single-AZ tasks
To the cloud	DB2 for LUW -> GaussDB Centralized	Not supported	Not supported
To the cloud	DB2 for LUW -> GaussDB Distributed	Not supported	Not supported
To the cloud	TiDB -> TaurusDB	Not supported	Not supported
To the cloud	Microsoft SQL Server -> GaussDB Centralized	Not supported	Not supported
To the cloud	Microsoft SQL Server -> GaussDB Distributed	Not supported	Not supported
To the cloud	Microsoft SQL Server -> Microsoft SQL Server	Not supported	Not supported
To the cloud	Microsoft SQL Server -> MySQL	Not supported	Not supported
To the cloud	Microsoft SQL Server -> TaurusDB	Not supported	Not supported
To the cloud	Microsoft SQL Server -> PostgreSQL	Not supported	Not supported
To the cloud	MongoDB -> DDS	Supported	Supported only for single-AZ tasks
To the cloud	MariaDB -> MariaDB	Supported	Supported only for single-AZ tasks

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
To the cloud	MariaDB -> MySQL	Supported	Supported only for single-AZ tasks
To the cloud	MariaDB -> TaurusDB	Supported	Supported only for single-AZ tasks
To the cloud	TaurusDB -> TaurusDB	Supported	Supported only for single-AZ tasks
To the cloud	Cassandra -> GeminiDB Cassandra	Not supported	Not supported
To the cloud	DynamoDB -> GeminiDB DynamoDB-Compatible API	Not supported	Not supported
From the cloud	MySQL -> MySQL	Supported	Supported only for single-AZ tasks
From the cloud	MySQL -> Kafka	Supported	Supported only for single-AZ tasks
From the cloud	MySQL -> Oracle	Supported	Supported only for single-AZ tasks
From the cloud	MySQL -> MariaDB	Not supported	Not supported
From the cloud	DDM -> MySQL	Supported	Not supported
From the cloud	DDM -> Oracle	Supported	Not supported
From the cloud	DDM -> Kafka	Supported	Not supported
From the cloud	DDS -> MongoDB	Supported	Not supported

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
From the cloud	DDS -> Kafka	Supported	Not supported
From the cloud	PostgreSQL -> PostgreSQL	Supported	Supported only for single-AZ tasks
From the cloud	PostgreSQL -> Kafka	Supported	Supported only for single-AZ tasks
From the cloud	GaussDB Centralized -> MySQL	Supported	Supported only for single-AZ tasks
From the cloud	GaussDB Centralized -> Oracle	Supported	Supported only for single-AZ tasks
From the cloud	GaussDB Centralized -> Kafka	Supported	Supported only for single-AZ tasks
From the cloud	GaussDB Centralized -> GaussDB Distributed	Supported	Supported only for single-AZ tasks
From the cloud	GaussDB Centralized -> GaussDB Centralized	Supported	Supported only for single-AZ tasks
From the cloud	GaussDB Distributed -> MySQL	Supported	Not supported
From the cloud	GaussDB Distributed -> Oracle	Supported	Not supported
From the cloud	GaussDB Distributed -> Kafka	Supported	Not supported
From the cloud	GaussDB Distributed -> GaussDB Distributed	Supported	Not supported
From the cloud	GaussDB Distributed -> GaussDB Centralized	Supported	Not supported

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
From the cloud	TaurusDB -> MySQL	Supported	Supported only for single-AZ tasks
From the cloud	TaurusDB -> Kafka	Supported	Supported only for single-AZ tasks
From the cloud	TaurusDB -> Oracle	Supported	Supported only for single-AZ tasks
From the cloud	MariaDB -> MariaDB	Supported	Supported only for single-AZ tasks
From the cloud	Microsoft SQL Server -> Kafka	Not supported	Not supported
Self-built -> Self-built	MySQL -> Kafka	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	MySQL -> GaussDB Centralized	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	MySQL -> GaussDB Distributed	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	Oracle -> Kafka	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	Oracle -> GaussDB Centralized	Supported	Supported only for single-AZ tasks

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
Self-built -> Self-built	Oracle -> GaussDB Distributed	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	PostgreSQL -> Kafka	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	GaussDB Centralized -> MySQL	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	GaussDB Centralized -> Oracle	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	GaussDB Centralized -> Kafka	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	GaussDB Centralized -> GaussDB Distributed	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	GaussDB Centralized -> GaussDB Centralized	Supported	Supported only for single-AZ tasks
Self-built -> Self-built	GaussDB Distributed -> MySQL	Supported	Not supported
Self-built -> Self-built	GaussDB Distributed -> Oracle	Supported	Not supported
Self-built -> Self-built	GaussDB Distributed -> Kafka	Supported	Not supported
Self-built -> Self-built	GaussDB Distributed -> GaussDB Centralized	Supported	Not supported

Synchronization Direction	Data Flow	Multiple Specifications	Specification Upgrade
Self-built -> Self-built	GaussDB Distributed -> GaussDB Distributed	Supported	Not supported
Self-built -> Self-built	DB2 for LUW -> GaussDB Centralized	Not supported	Not supported
Self-built -> Self-built	DB2 for LUW -> GaussDB Distributed	Not supported	Not supported
Self-built -> Self-built	Microsoft SQL Server -> Kafka	Not supported	Not supported

7.2 Real-Time Disaster Recovery

Precautions

The performance indicators provided in this section are for reference only. The actual environment is affected by factors such as the performance of the source or destination database, network bandwidth, data model, and service model.

Specification Description

Based on the maximum performance of data flows, there are four types of specifications: micro, small, medium, and large. [Table 7-5](#) lists the maximum performance of each specification.

Table 7-5 Maximum performance

Specifications	Reference Value of Maximum Performance (Rows/Second)
Micro	300
Small	3,000
Medium	7,500
Large	10,000

 NOTE

- The performance of each specification is affected by factors such as the networks, service and DR database performance, and latency. The values in the table are for reference only.
- DRS provides specifications of different performance which is measured by DR initialization (no flow control) and DR performance.
- The maximum performance (rows/second) is measured by the number of DR rows per second, including every row of data changes caused by DML statements (INSERT, DELETE, and UPDATE). You can [view the destination database write frequency \(apply_rows_rate\) monitoring metric](#) on Cloud Eye.
- DRS allows you to upgrade specifications only for DR tasks from MySQL to MySQL, MySQL to TaurusDB, and TaurusDB to TaurusDB. Task specifications cannot be downgraded. For details, see [Changing Specifications](#).
- If you want to compare values for a DRS task, select large specifications when creating the DRS task.

Multiple Specifications

DRS real-time DR allows you to select specifications for some specified data flow tasks.

Table 7-6 Data types that support multiple specifications

DR Direction	Data Flow	Multiple Specifications	Specification Upgrade
Current cloud as standby	MySQL -> MySQL	Supported	Yes
Current cloud as active	MySQL -> MySQL	Supported	Yes
Current cloud as standby	MySQL -> TaurusDB	Supported	Yes
Current cloud as standby	DDM -> DDM	Unsupported	No
Current cloud as active	DDM -> DDM	Unsupported	No
Current cloud as standby	TaurusDB -> TaurusDB	Supported	Yes
Current cloud as active	TaurusDB -> TaurusDB	Supported	Yes

DR Direction	Data Flow	Multiple Specifications	Specification Upgrade
Dual-active DR	MySQL -> MySQL	Supported	Yes
Dual-active DR	TaurusDB -> TaurusDB	Supported	No

8 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.

8.1 MySQL->PostgreSQL

Table 8-1 Data type mapping

Data Type (MySQL)	Data Type (PostgreSQL)	Whether to Support Mapping
BIGINT	NUMERIC BIGINT	Yes
BINARY	BYTEA	Yes
BIT	BIT	Yes
BLOB	BYTEA	Yes
BOOLEAN	BOOL	Yes
CHAR	CHAR	Yes
DATE	DATE	Yes
DATETIME	TIMESTAMP	Yes
DECIMAL	NUMERIC	Yes
DOUBLE	FLOAT8	Yes
ENUM	VARCHAR	Yes
FLOAT	FLOAT4 FLOAT8	Yes
INT	INTEGER	Yes
LONGBLOB	BYTEA	Yes

Data Type (MySQL)	Data Type (PostgreSQL)	Whether to Support Mapping
LONGTEXT	TEXT	Yes
MEDIUMBLOB	BYTEA	Yes
MEDIUMINT	INT	Yes
SET	VARCHAR	Yes
SMALLINT	INT SMALLINT	Yes
TEXT	TEXT	Yes
TIME	TIME	Yes
TIMESTAMP	TIMESTAMP	Yes
TINYBLOB	BYTEA	Yes
TINYINT	SMALLINT	Yes
TINYTEXT	TEXT	Yes
VARBINARY	BYTEA	Yes
VARCHAR	VARCHAR	Yes
YEAR	SMALLINT	Yes
JSON	JSON	Yes
GEOMETRY	-	No

 NOTE

- DATE values supported by MySQL range from '1000-01-01' to '9999-12-31'.
DATETIME values supported by MySQL range from '1000-01-01 00:00:00' to '9999-12-31 23:59:59'.
TIMESTAMP values supported by MySQL range from '1970-01-01 00:00:01' UTC to '2038-01-19 03:14:07' UTC.
For details, see [official MySQL documentation](#).
For PostgreSQL, 0000-00-00 is an invalid date and will be converted to 1970-01-01 by DRS. For example, '0000-00-00' of the DATE type in MySQL is converted to '1970-01-01' by DRS, and '1000-00-31 23:59:59' of the DATETIME or TIMESTAMP type in MySQL is converted to '1970-01-01 00:00:00' by DRS.
- TIME values supported by MySQL range from '-838:59:59' to '838:59:59'. For details, see the [official MySQL documentation](#). For PostgreSQL, the minimum value of the TIME type is 00:00:00 and the maximum value is 24:00:00. In MySQL, if a value of the TIME type is less than 00:00:00 or greater than 24:00:00, DRS will convert it to 00:00:00.
- YEAR value ranges supported by MySQL are 1901 to 2155 and 0000. For details, see [official MySQL documentation](#). PostgreSQL does not have the YEAR type, so DRS will convert the YEAR type of MySQL to the SMALLINT type.
- For MySQL databases, '0000' of the DATE type will be converted to 0 by DRS.
- If the data type of a column is INT and the column contains the AUTO_INCREMENT attribute, DRS converts the data type of the column to SERIAL during synchronization.

8.2 MySQL->GaussDB

MySQL -> GaussDB B-, MySQL-, and PostgreSQL-Compatible Mode

Table 8-2 Data type mapping

Data Type (MySQL)	Data Type (GaussDB)	Whether to Support Mapping
CHAR	CHARACTER	Yes. If a column of this type in the source database contains characters that occupy more than one byte, increase the length of the column in the destination database.
VARCHAR	CHARACTER VARYING()	Yes. If a column of this type in the source database contains characters that occupy more than one byte, increase the length of the column in the destination database.
BINARY	BYTEA	Yes

Data Type (MySQL)	Data Type (GaussDB)	Whether to Support Mapping
VARBINARY	BYTEA	Yes
TINYBLOB	BYTEA	Yes
BLOB	BYTEA	Yes
MEDIUMBLOB	BYTEA	Yes
LOB	BYTEA	Yes
TINYTEXT	TEXT	Yes
TEXT	TEXT	Yes
MEDIUMTEXT	TEXT	Yes
LONGTEXT	CLOB	Yes
ENUM	VARCHAR	Yes
SET	VARCHAR	Yes
TINYINT	SMALLINT	Yes
SMALLINT	SMALLINT	Yes
MEDIUMINT	INT	Yes
INT	INT	Yes
BIGINT	BIGINT	Yes
FLOAT	REAL/DOUBLE PRECISION	Yes
DOUBLE	DOUBLE PRECISION	Yes
DATE	DATE	Yes
DATETIME	TIMESTAMP WITHOUT TIME ZONE	Yes
TIMESTAMP	TIMESTAMP WITH TIME ZONE	Yes
TIME	TIME WITHOUT TIME ZONE	Yes
BIT	BIT	Yes
JSON	JSON	Yes, but JSON containing BIT data is not supported.
DECIMAL	NUMERIC	Yes
NUMERIC	NUMERIC	Yes

Data Type (MySQL)	Data Type (GaussDB)	Whether to Support Mapping
YEAR	SMALLINT	Yes
BOOLEAN	SMALLINT	Yes

 **NOTE**

- DATE values supported by MySQL range from '1000-01-01' to '9999-12-31'.
DATETIME values supported by MySQL range from '1000-01-01 00:00:00' to '9999-12-31 23:59:59'.
TIMESTAMP values supported by MySQL range from '1970-01-01 00:00:01' UTC to '2038-01-19 03:14:07' UTC.
For details, see [official MySQL documentation](#).
- YEAR value ranges supported by MySQL are 1901 to 2155 and 0000. For details, see [official MySQL documentation](#). GaussDB does not have the YEAR type, so DRS will convert the YEAR type of MySQL to the SMALLINT type.
- For MySQL databases, '0000' of the DATE type will be converted to 0 by DRS.
- If the data type of a column is INT and the column contains the AUTO_INCREMENT attribute, DRS converts the data type of the column to SERIAL during synchronization.
- Auto-increment columns cannot be migrated in PostgreSQL compatibility mode.

MySQL -> GaussDB M-Compatible Mode

Table 8-3 Data type mapping

Data Type (MySQL)	Data Type (GaussDB)	Whether to Support Mapping
BOOL	BOOL	Yes
BOOLEAN	BOOLEAN	Yes
TINYINT	TINYINT	Yes
SMALLINT	SMALLINT	Yes
MEDIUMINT	MEDIUMINT	Yes
INT	INT	Yes
INTEGER	INTEGER	Yes
BIGINT	BIGINT	Yes
DECIMAL	DECIMAL	Yes
NUMERIC	DECIMAL	Yes
DEC	DEC	Yes
FIXED	DECIMAL	Yes

Data Type (MySQL)	Data Type (GaussDB)	Whether to Support Mapping
FLOAT	FLOAT	Yes
DOUBLE	DOUBLE	Yes
DOUBLE PRECISION	DOUBLE	Yes
REAL	DOUBLE	Yes
DATE	DATE	Yes
DATETIME	DATETIME	Yes
TIMESTAMP	TIMESTAMP	Yes
TIME	TIME	Yes
YEAR	YEAR	Yes
CHAR	CHAR	Yes
VARCHAR	VARCHAR	Yes
TINYTEXT	TINYTEXT	Yes
TEXT	TEXT	Yes
MEDIUMTEXT	MEDIUMTEXT	Yes
LONGTEXT	LONGTEXT	Yes
ENUM('value1','value2',...)	VARCHAR+CHECK	Yes
SET('value1','value2',...)	VARCHAR	Yes
BINARY	BINARY	Yes
VARBINARY	VARBINARY	Yes
TINYBLOB	TINYBLOB	Yes
BLOB	BLOB	Yes
MEDIUMBLOB	MEDIUMBLOB	Yes
LOB	LOB	Yes
BIT	BIT	Yes
JSON	JSON	Yes

8.3 DDM->GaussDB

DDM -> GaussDB B-, MySQL-, and PostgreSQL-Compatible Mode

Table 8-4 Data type mapping

Data Type (DDM)	Data Type (GaussDB)	Whether to Support Mapping
CHAR	CHARACTER	Yes. If a column of this type in the source database contains characters that occupy more than one byte, increase the length of the column in the destination database.
VARCHAR	CHARACTER VARYING()	Yes. If a column of this type in the source database contains characters that occupy more than one byte, increase the length of the column in the destination database.
BINARY	BYTEA	Yes
VARBINARY	BYTEA	Yes
TINYBLOB	BYTEA	Yes
BLOB	BYTEA	Yes
MEDIUMBLOB	BYTEA	Yes
LOB	BYTEA	Yes
TINYTEXT	TEXT	Yes
TEXT	TEXT	Yes
MEDIUMTEXT	TEXT	Yes
LONGTEXT	CLOB	Yes
ENUM	VARCHAR	Yes
SET	VARCHAR	Yes
TINYINT	SMALLINT	Yes
SMALLINT	SMALLINT	Yes

Data Type (DDM)	Data Type (GaussDB)	Whether to Support Mapping
MEDIUMINT	INT	Yes
INT	INT	Yes
BIGINT	BIGINT	Yes
FLOAT	REAL/DOUBLE PRECISION	Yes
DOUBLE	DOUBLE PRECISION	Yes
DATE	DATE	Yes
DATETIME	TIMESTAMP WITHOUT TIME ZONE	Yes
TIMESTAMP	TIMESTAMP WITH TIME ZONE	Yes
TIME	TIME WITHOUT TIME ZONE	Yes
BIT	BIT	Yes
JSON	JSON	Yes, but JSON containing BIT data is not supported.
DECIMAL	NUMERIC	Yes
NUMERIC	NUMERIC	Yes
YEAR	SMALLINT	Yes
BOOLEAN	SMALLINT	Yes

NOTE

- DATE values supported by DDM range from '1000-01-01' to '9999-12-31'.
DATETIME values supported by DDM range from '1000-01-01 00:00:00' to '9999-12-31 23:59:59'.
TIMESTAMP values supported by DDM range from '1970-01-01 00:00:01' UTC to '2038-01-19 03:14:07' UTC.
For details, see [official MySQL documentation](#).
- YEAR value ranges supported by DDM are 1901 to 2155 and 0000. For details, see [official MySQL documentation](#). GaussDB does not have the YEAR type, so DRS will convert the YEAR type of DDM to the SMALLINT type.
- For DDM databases, '0000' of the DATE type will be converted to 0 by DRS.
- If the data type of a column is INT and the column contains the AUTO_INCREMENT attribute, DRS converts the data type of the column to SERIAL during synchronization.
- Auto-increment columns cannot be migrated in PostgreSQL compatibility mode.

DDM -> GaussDB M-Compatible Mode

Table 8-5 Data type mapping

Data Type (DDM)	Data Type (GaussDB)	Whether to Support Mapping
BOOL	BOOL	Yes
BOOLEAN	BOOLEAN	Yes
TINYINT	TINYINT	Yes
SMALLINT	SMALLINT	Yes
MEDIUMINT	MEDIUMINT	Yes
INT	INT	Yes
INTEGER	INTEGER	Yes
BIGINT	BIGINT	Yes
DECIMAL	DECIMAL	Yes
NUMERIC	DECIMAL	Yes
DEC	DEC	Yes
FIXED	DECIMAL	Yes
FLOAT	FLOAT	Yes
DOUBLE	DOUBLE	Yes
DOUBLE PRECISION	DOUBLE	Yes
REAL	DOUBLE	Yes
DATE	DATE	Yes
DATETIME	DATETIME	Yes
TIMESTAMP	TIMESTAMP	Yes
TIME	TIME	Yes
YEAR	YEAR	Yes
CHAR	CHAR	Yes
VARCHAR	VARCHAR	Yes
TINYTEXT	TINYTEXT	Yes
TEXT	TEXT	Yes
MEDIUMTEXT	MEDIUMTEXT	Yes
LONGTEXT	LONGTEXT	Yes

Data Type (DDM)	Data Type (GaussDB)	Whether to Support Mapping
ENUM('value1','value2',...)	VARCHAR	Yes
SET('value1','value2',...)	VARCHAR+CHECK	Yes
BINARY	BINARY	Yes
VARBINARY	VARBINARY	Yes
TINYBLOB	TINYBLOB	Yes
BLOB	BLOB	Yes
MEDIUMBLOB	MEDIUMBLOB	Yes
LONGBLOB	LONGBLOB	Yes
BIT	BIT	Yes
JSON	JSON	Yes

8.4 MySQL->Oracle

Table 8-6 Data type mapping

Data Type (MySQL)	Data Type (Oracle)	Whether to Support Mapping
ENUM	VARCHAR2	Yes
SET	VARCHAR2	Yes
VARCHAR	VARCHAR2	Yes
NUMERIC	NUMBER	Yes
FLOAT	BINARY_FLOAT	Yes
TIMESTAMP	TIMESTAMP WITH TIME ZONE	Yes
DATETIME	TIMESTAMP	Yes
DATE	DATE	Yes
TIME	INTERVAL DAY TO SECOND	Yes
YEAR	VARCHAR2	Yes
BIT	RAW	Yes
CLOB	CLOB	Yes

Data Type (MySQL)	Data Type (Oracle)	Whether to Support Mapping
GEOMETRY	-	No
VARBINARY	BLOB	Yes
BINARY	RAW	Yes
DOUBLE	BINARY_DOUBLE	Yes
DECIMAL	NUMBER	Yes
INT	NUMBER	Yes
TINYINT	NUMBER	Yes
SMALLINT	NUMBER	Yes
MEDIUMINT	NUMBER	Yes
BIGINT	NUMBER	Yes
BLOB	BLOB	Yes
LOB	BLOB	Yes
MEDIUMBLOB	BLOB	Yes
CHAR	CHAR	Yes
TEXT	CLOB	Yes
JSON	CLOB	Yes

8.5 Oracle->MySQL

Table 8-7 Data type mapping

Data Type (Oracle)	Condition	Data Type (MySQL)	Whether to Support Mapping
CHAR	length<=255	CHAR	Yes
CHAR	length>255	VARCHAR	Yes
VARCHAR	Length (row size) ≤ 65536	VARCHAR	Yes
VARCHAR	Length (row size) > 65536	TEXT	Yes
VARCHAR2	-	VARCHAR2	Yes
NCHAR	length<=255	NCHAR	Yes

Data Type (Oracle)	Condition	Data Type (MySQL)	Whether to Support Mapping
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	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	DATETIME	Yes
TIMESTAMP WITH TIME ZONE	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	TIMESTAMP	Yes

Data Type (Oracle)	Condition	Data Type (MySQL)	Whether to Support Mapping
TIMESTAMP WITH LOCAL TIME ZONE	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	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	-	LONGTEXT	Yes
BFILE	-	-	No
SDO_GEOMETRY	-	-	No

8.6 Oracle->TaurusDB

Table 8-8 Data type mapping

Data Type (Oracle)	Condition	Data Type (TaurusDB)	Whether to Support Mapping
CHAR	length<=255	CHAR	Yes
CHAR	length>255	VARCHAR	Yes
VARCHAR	Size<=65536	VARCHAR	Yes
VARCHAR	Size>65536	TEXT	Yes

Data Type (Oracle)	Condition	Data Type (TaurusDB)	Whether to Support Mapping
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	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	DATETIME	Yes
TIMESTAMP WITH TIME ZONE	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	TIMESTAMP	Yes

Data Type (Oracle)	Condition	Data Type (TaurusDB)	Whether to Support Mapping
TIMESTAMP WITH LOCAL TIME ZONE	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	TIMESTAMP	Yes
INTERVAL	Incremental	VARCHAR(30)	No
INTERVAL	Full; 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	-	LONGTEXT	Yes
BFILE	-	-	No
SDO_GEOMETRY	-	-	No

8.7 Oracle->GaussDB

Table 8-9 Data type mapping

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
CHAR	CHARACTER	Supported	Supported	Supported. The spaces before and after the character are ignored.	Supported. The spaces before and after the character are ignored.	-
VARCHAR	CHARACTER VARYING	Supported	Supported	Supported	Supported	-
VARCHAR2	CHARACTER VARYING	Supported	Supported	Supported	Supported	-
NCHAR	CHARACTER	Supported	Supported	Supported. The spaces before and after the character are ignored.	Supported. The spaces before and after the character are ignored.	-
NVARCHAR2	NVARCHAR2	Supported	Supported	Supported	Supported	-
NUMBER	NUMERIC	Supported	Supported	Supported	Supported	-
NUMBER (6,3)	NUMERIC(6,3)	Supported	Supported	Supported	Supported	-
NUMBER (6,0)	INTEGER	Supported	Supported	Supported	Supported	-

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
NUMBER (3)	SMALLINT	Supported	Supported	Supported	Supported	-
NUMBER (6,-2)	INTEGER	Supported	Supported	Supported	Supported	-
BINARY_FLOAT	REAL	Supported	Supported	Not supported	Supported	The precision ranges of the source and destination databases are different, which may cause precision loss.
BINARY_DOUBLE	DOUBLE PRECISION	Supported	Supported	Not supported	Supported	The precision ranges of the source and destination databases are different, which may cause precision loss.

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
FLOAT	DOUBLE PRECISION	Supported	Supported	Not supported	Supported	The precision ranges of the source and destination databases are different, which may cause precision loss.
INT	NUMERIC	Supported	Supported	Supported	Supported	-
INTEGER	NUMERIC	Supported	Supported	Supported	Supported	-

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
DATE	TIMESTAMP(0) WITHOUT TIME ZONE	Supported	Supported	Not supported	Supported	If a table with date type is created in the destination database, the data type precision range in the source database is different from that in the destination database, causing precision loss. Therefore, comparison is not supported.

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
TIMESTAMP	TIMESTAMP(6) WITHOUT TIME ZONE	Supported	Supported	Not supported	The value is accurate to six decimal places.	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.
TIMESTAMP_TZ	TIMESTAMP(6) WITH TIME ZONE	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Filter this column.	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
TIMESTAMP_LTZ	TIMESTAMP(6) WITH TIME ZONE	Not supported (The destination database does not support creating tables using the primary key.)	Supported	Not supported	Filter this column.	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.
INTERVAL_YM	INTERVAL YEAR TO MONTH	Supported	Supported	Not supported	Not supported	Incremental synchronization does not support this type.
INTERVAL_DS	INTERVAL DAY TO SECOND	Supported	Supported	Not supported	Not supported	Incremental synchronization does not support this type. The maximum precision supported by the source database is 6.

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
BLOB	BLOB	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Supported	You can choose to filter data or compare the length, hash, and content. During hash comparison, Oracle uses the hash function in the DBMS_CRYPTO package to obtain the LOB hash value. To use the DBMS_CRYPTO package, the SYSDBA needs to grant permissions to users. Reference statement: GRANT EXECUTE ON DBMS_CRYPTO

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
CLOB	CLOB	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Supported	You can choose to filter data or compare the length, hash, and content. During hash comparison, Oracle uses the hash function in the DBMS_CRYPTO package to obtain the LOB hash value. To use the DBMS_CRYPTO package, the SYSDBA needs to grant permissions to users. Reference statement: GRANT EXECUTE ON DBMS_CRYPTO

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
NCLOB	TEXT	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Supported	You can choose to filter data or compare the length, hash, and content. During hash comparison, Oracle uses the hash function in the DBMS_CRYPTO package to obtain the LOB hash value. To use the DBMS_CRYPTO package, the SYSDBA needs to grant permissions to users. Reference statement: GRANT EXECUTE ON DBMS_CRYPTO

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
LONG	TEXT	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Filter this column.	-
LONG_RAW	BYTEA	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Filter this column.	-
RAW	RAW	Supported	Supported	Not supported	Supported	-
ROWID	CHARACTER(18)	Supported	Supported	Supported	Supported	-

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
BFILE	-	Not supported	Not supported	Not supported	Not supported	Restrictions on the source database: The BFILE type is not supported.
XMLTYPE	XMLTYPE	Not supported (The source database does not support creating tables using the primary key.)	Supported	Not supported	Not supported	-
UROWID	-	Not supported	Not supported	Not supported	Not supported	Full and incremental synchronizations are not supported.

Source Data Type	Destination Data Type	Synchronization (Source Data Type as Primary Key)	Synchronization (Source Data Type as Non-Primary Key)	Comparison (Source Data Type as Primary Key)	Comparison (Source Data Type as Non-Primary Key)	Remarks
SDO_GEOMETRY	-	Not supported	Not supported	Not supported	Not supported	Restrictions on the source database: The SDO_GEOMETRY type is not supported.
NUMBER(*, 0)	NUMERIC	Supported	Supported	Supported	Supported	-

8.8 Oracle->DDM

Table 8-10 Data type mapping

Data Type (Oracle)	Condition	Data Type (DDM)	Whether to Support Mapping
CHAR(n)	n<=255	CHAR(n)	Yes
CHAR(n)	n>255	VARCHAR(n)	Yes
VARCHAR(Size)	Length (row size) ≤ 65535	VARCHAR(n)	Yes
VARCHAR(Size)	Length (row size) > 65535	TEXT	Yes
VARCHAR2(n)	-	VARCHAR(n)	Yes
NCHAR(n)	n<=255	NCHAR(n)	Yes
NCHAR(n)	n>255	NVARCHAR(n)	Yes
NVARCHAR2(n)	-	NVARCHAR(n)	Yes
NUMBER(p,s)	s>0	NUMBER(p,s)	Yes
NUMBER(p,s)	s<=0	NUMBER(p-s,0)	Yes

Data Type (Oracle)	Condition	Data Type (DDM)	Whether to Support Mapping
BINARY_FLOAT	-	FLOAT	Yes
BINARY_DOUBLE	-	DOUBLE	Yes
FLOAT(b)	b<=99	DECIMAL(b*0.30103*2, b*0.30103)	Yes
FLOAT(b)	b>99	DOUBLE	Yes
DATE	-	DATETIME	Yes
TIMESTAMP	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	TIMESTAMP	Yes
TIMESTAMP WITH LOCAL TIME ZONE	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	TIMESTAMP	Yes

Data Type (Oracle)	Condition	Data Type (DDM)	Whether to Support Mapping
TIMESTAMP WITH TIME ZONE	If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.	TIMESTAMP	Yes
INTERVAL	Incremental	VARCHAR(30)	No
INTERVAL	Full; 6 digit precision	VARCHAR(30)	Yes
BLOB	-	LOB	Yes
CLOB	-	LONGTEXT	Yes
NCLOB	-	LONGTEXT	Yes
LONG	-	LONGTEXT	Yes
LONG_RAW	-	LOB	Yes
RAW	-	VARBINARY	Yes
ROWID	-	VARCHAR(18)	Yes
XMLTYPE	-	LONGTEXT	Yes

8.9 Oracle->PostgreSQL

Oracle -> PostgreSQL Community Edition

Table 8-11 Data type mapping

Data Type (Oracle)	Data Type (PostgreSQL Community Edition)	Whether to Support Mapping
CHAR	CHAR	Yes
VARCHAR	VARCHAR	Yes
VARCHAR2	VARCHAR	Yes

Data Type (Oracle)	Data Type (PostgreSQL Community Edition)	Whether to Support Mapping
NCHAR	NCHAR	Yes
NVARCHAR2	VARCHAR	Yes
NUMBER	NUMBER	Yes
BINARY_FLOAT	REAL	Yes
BINARY_DOUBLE	DOUBLE	Yes
FLOAT	FLOAT	Yes
DATE	TIMESTAMP	Yes
TIMESTAMP	TIMESTAMP	Yes. If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.
TIMESTAMP WITH TIME ZONE	TIMESTAMPTZ	Yes. If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.
TIMESTAMP WITH LOCAL TIME ZONE	TIMESTAMPTZ	Yes. If the precision of the source database exceeds 6 digits, the precision will be reduced to 6 digits due to the maximum precision allowed in the destination database.
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	CHARACTER(18)	Yes
UROWID	-	No
XMLTYPE	TEXT	Yes
BFILE	-	No
SDO_GEOMETRY	-	No

8.10 TaurusDB->Oracle

Table 8-12 Data type mapping

Data Type (TaurusDB)	Data Type (Oracle)	Whether to Support Mapping
ENUM	VARCHAR2	Yes
SET	VARCHAR2	Yes
VARCHAR	VARCHAR2	Yes
NUMERIC	NUMBER	Yes
FLOAT	BINARY_FLOAT	Yes
TIMESTAMP	TIMESTAMP WITH TIME ZONE	Yes
DATETIME	TIMESTAMP	Yes
DATE	DATE	Yes
TIME	INTERVAL DAY TO SECOND	Yes
YEAR	VARCHAR2	Yes
BIT	RAW	Yes
CLOB	CLOB	Yes
VARBINARY	BLOB	Yes
BINARY	RAW	Yes

Data Type (TaurusDB)	Data Type (Oracle)	Whether to Support Mapping
DOUBLE	BINARY_DOUBLE	Yes
DECIMAL	NUMBER	Yes
INT	NUMBER	Yes
TINYINT	NUMBER	Yes
SMALLINT	NUMBER	Yes
MEDIUMINT	NUMBER	Yes
BIGINT	NUMBER	Yes
BLOB	BLOB	Yes
LOB	BLOB	Yes
MEDIUMBLOB	BLOB	Yes
CHAR	CHAR	Yes
TEXT	CLOB	Yes
JSON	CLOB	Yes
GEOMETRY	-	No

8.11 GaussDB->MySQL

GaussDB B or MySQL-Compatible Mode -> MySQL

Table 8-13 Data type mapping

Data Type (GaussDB)	Data Type (MySQL)	Whether to Support Mapping
CHARACTER	CHAR	Yes
CHARACTER VARYING	VARCHAR	Yes
BYTEA/BLOB/RAW	BINARY/VARBINARY/ TINYBLOB /BLOB/ MEDIUMBLOB/ LOB	Yes
TEXT/CLOB	TINYTEXT/ MEDIUMTEXT / LONGTEXT/ TEXT	Yes
TINYINT	SMALLINT	Yes

Data Type (GaussDB)	Data Type (MySQL)	Whether to Support Mapping
SMALLINT	SMALLINT	Yes
INTEGER	INT	Yes
BIGINT	BIGINT	Yes
FLOAT4	FLOAT	Yes
FLOAT8/DOUBLE PRECISION	DOUBLE	Yes
DATE	DATE	Yes
TIMESTAMP WITHOUT TIME ZONE	DATETIME	Yes
SMALLDATETIME	DATETIME	Yes
TIMESTAMP WITH TIME ZONE	TIMESTAMP	Yes
TIME WITH TIME ZONE	TIME	Yes
TIME WITHOUT TIME ZONE	TIME	Yes
BIT	BIT	Yes
MONEY	VARCHAR	Yes
BOOLEAN	BOOLEAN	Yes
NUMBER	DECIMAL	Yes
NUMERIC	DECIMAL	Yes
DECIMAL	DECIMAL	Yes
TINYINT UNSIGNED	TINYINT UNSIGNED	Yes
SMALLINT UNSIGNED	SMALLINT UNSIGNED	Yes
INTEGER UNSIGNED	INTEGER UNSIGNED	Yes
BIGINT UNSIGNED	BIGINT UNSIGNED	Yes

GaussDB M-Compatible Mode -> MySQL

Table 8-14 Data type mapping

Data Type (GaussDB)	Data Type (MySQL)	Whether to Support Mapping
BOOL	BOOL	Yes
BOOLEAN	BOOLEAN	Yes
TINYINT	TINYINT	Yes
SMALLINT	SMALLINT	Yes
MEDIUMINT	MEDIUMINT	Yes
INT	INT	Yes
INTEGER	INTEGER	Yes
BIGINT	BIGINT	Yes
TINYINT UNSIGNED	TINYINT UNSIGNED	Yes
SMALLINT UNSIGNED	SMALLINT UNSIGNED	Yes
MEDIUMINT UNSIGNED	MEDIUMINT UNSIGNED	Yes
INTEGER UNSIGNED	INTEGER UNSIGNED	Yes
BIGINT UNSIGNED	BIGINT UNSIGNED	Yes
DECIMAL	DECIMAL	Yes
NUMERIC	DECIMAL	Yes
DEC	DECIMAL	Yes
FIXED	DECIMAL	Yes
FLOAT	FLOAT	Yes
DOUBLE	DOUBLE	Yes
DOUBLE PRECISION	DOUBLE	Yes
REAL	DOUBLE	Yes
DATE	DATE	Yes
DATETIME	DATETIME	Yes
TIMESTAMP	TIMESTAMP	Yes
TIME	TIME	Yes
YEAR	YEAR	Yes
CHAR	VARCHAR	Yes

Data Type (GaussDB)	Data Type (MySQL)	Whether to Support Mapping
VARCHAR	VARCHAR/ENUM/SET	Yes
TINYTEXT	TINYTEXT	Yes
TEXT	TEXT/JSON	Yes
MEDIUMTEXT	MEDIUMTEXT	Yes
LONGTEXT	LONGTEXT	Yes
BINARY	BINARY	Yes
VARBINARY	VARBINARY	Yes
TINYBLOB	TINYBLOB	Yes
BLOB	BLOB	Yes
MEDIUMBLOB	MEDIUMBLOB	Yes
LOB	LOB	Yes
BIT	BIT	Yes
JSON	JSON	Yes

8.12 GaussDB->PostgreSQL

Table 8-15 Data type mapping

Data Type (GaussDB)	Data Type (PostgreSQL)	Whether to Support Mapping
BOOLEAN	BOOLEAN	Yes
SMALLINT	SMALLINT	Yes
TINYINT	SMALLINT	Yes
INTEGER	INTEGER	Yes
BIGINT	BIGINT	Yes
NUMERIC	NUMERIC	Yes
NUMBER	NUMERIC	Yes
REAL	REAL	Yes
DOUBLE PRECISION	DOUBLE PRECISION	Yes
SMALLSERIAL	SMALLSERIAL	Yes

Data Type (GaussDB)	Data Type (PostgreSQL)	Whether to Support Mapping
SERIAL	SERIAL	Yes
BIGSERIAL	BIGSERIAL	Yes
LARGESERIAL	BIGSERIAL	Yes. Data overflow may occur because the value range of the LARGESERIAL data type supported by GaussDB is greater than the maximum value of BIGSERIAL for the SERIAL data type supported by PostgreSQL.
CHARACTER(n)	CHARACTER(n)	Yes
CHAR(n)	CHAR(n)	Yes
NCHAR(n)	CHAR(n)	Yes
CHARACTER VARYING(n)	CHARACTER VARYING(n)	Yes
VARCHAR(n)	VARCHAR(n)	Yes
VARCHAR2(n)	VARCHAR(n)	Yes
NVARCHAR2(n)	VARCHAR(n)	Yes
NVARCHAR(n)	VARCHAR(n)	Yes
TEXT	TEXT	Yes
CLOB	TEXT	Yes
DATE	DATE	Yes
TIME [(p)] [WITHOUT TIME ZONE]	TIME [(p)] [WITHOUT TIME ZONE]	Yes
TIME [(p)] [WITH TIME ZONE]	TIME [(p)] [WITH TIME ZONE]	Yes
TIMESTAMP[(p)] [WITHOUT TIME ZONE]	TIMESTAMP[(p)] [WITHOUT TIME ZONE]	Yes
TIMESTAMP[(p)] [WITH TIME ZONE]	TIMESTAMP[(p)] [WITH TIME ZONE]	Yes
INTERVAL DAY (l) TO SECOND (p)	INTERVAL[fields] [(p)]	Yes
INTERVAL[fields] [(p)]	INTERVAL[fields] [(p)]	Yes
BLOB	BYTEA	Yes

Data Type (GaussDB)	Data Type (PostgreSQL)	Whether to Support Mapping
RAW	BYTEA	Yes. RAW data supported by GaussDB, a hexadecimal string, will be written into the BYTEA column supported by PostgreSQL. The length in bytes of a RAW value in GaussDB is twice that of a BYTEA value in PostgreSQL. GaussDB and PostgreSQL use the same bytecode and hexadecimal notation.
BYTEA	BYTEA	Yes
BIT(n)	BIT(n)	Yes
BIT VARYING(n)	BIT VARYING(n)	Yes
POINT	POINT	Yes
LSEG	LSEG	Yes
BOX	BOX	Yes
PATH	PATH	Yes
POLYGON	POLYGON	Yes
CIRCLE	CIRCLE	Yes
CIDR	CIDR	Yes
INET	INET	Yes
MACADDR	MACADDR	Yes
TSVECTOR	TSVECTOR	Yes
TSQUERY	TSQUERY	Yes
UUID	UUID	Yes
INTEGER[]	INTEGER[]	Yes
TEXT[][]	TEXT[][]	Yes
JSON	JSON	Yes
MONEY	MONEY	Yes
TINYINT UNSIGNED	SMALLINT	Yes
SMALLINT UNSIGNED	INT	Yes

Data Type (GaussDB)	Data Type (PostgreSQL)	Whether to Support Mapping
INTEGER UNSIGNED	BIGINT	Yes
BIGINT UNSIGNED	NUMERIC	Yes
SMALLDATETIME	-	No
RELTIME	-	No
ABSTIME	-	No
TID	-	No
XID	-	No
CID	-	No
OID	-	No

8.13 GaussDB->Oracle

Table 8-16 Data type mapping

Data Type (GaussDB)	Data Type (Oracle)	Whether to Support Mapping
TINYINT	NUMBER	Yes
SMALLINT	NUMBER	Yes
INTEGER	NUMBER	Yes
BIGINT	NUMBER	Yes
NUMBER	NUMBER	Yes
NUMERIC	NUMBER	Yes
REAL	BINARY_FLOAT	Yes
DOUBLE PRECISION	BINARY_DOUBLE	Yes
DATE	DATE	Yes
BOOLEAN	CHAR(1)	Yes
CHARACTER	CHAR	Yes
CHARACTER VARYING	VARCHAR2	Yes
NVARCHAR2	NVARCHAR2	Yes
TEXT	CLOB	Yes
BLOB	BLOB	Yes

Data Type (GaussDB)	Data Type (Oracle)	Whether to Support Mapping
BYTEA	BLOB	Yes
TIMESTAMP WITHOUT TIME ZONE	TIMESTAMP	Yes
TIMESTAMP WITH TIME ZONE	TIMESTAMP WITH TIME ZONE	Yes
TIME WITHOUT TIME ZONE	VARCHAR2(32)	Yes
TIME WITH TIME ZONE	VARCHAR2(32)	Yes
CLOB	CLOB	Yes
RAW	RAW	Yes
MONEY	VARCHAR2	Yes
TINYINT UNSIGNED	NUMBER	Yes
SMALLINT UNSIGNED	NUMBER	Yes
INTEGER UNSIGNED	NUMBER	Yes
BIGINT UNSIGNED	NUMBER	Yes

8.14 DB2 for LUW->GaussDB

Table 8-17 Data type mapping

Data Type (DB2 for LUW)	Data Type (GaussDB)	Whether to Support Mapping
CHARACTER	CHARACTER	Yes
VARCHAR	VARCHAR	Yes
LONG VARCHAR	TEXT	Yes
GRAPHIC	NVARCHAR2	Yes
VARGRAPHIC	NVARCHAR2	Yes
LONG VARGRAPHIC	TEXT	Yes
CLOB	CLOB	Yes
DBCLOB	TEXT	Yes
BLOB	BLOB	Yes
BINARY	RAW	Yes

Data Type (DB2 for LUW)	Data Type (GaussDB)	Whether to Support Mapping
VARBINARY	RAW	Yes
REAL	REAL	Yes
DOUBLE	FLOAT8	Yes
SMALLINT	SMALLINT	Yes
INTEGER	INTEGER	Yes
BIGINT	BIGINT	Yes
DECIMAL	DECIMAL	Yes
DECFLOAT	NUMERIC	Yes
DATE	DATE	Yes
TIME	TIME(0) WITHOUT TIME ZONE	Yes
TIMESTAMP	TIMESTAMP(6) WITHOUT TIME ZONE	Yes
XML	TEXT	Yes
BOOLEAN	BOOLEAN	Yes
DB2SECURITYLABEL	VARCHAR (128)	Yes

8.15 PostgreSQL->GaussDB

Table 8-18 Data type mapping

Data Type (PostgreSQL)	Data Type (GaussDB)	Whether to Support Mapping
SMALLINT	SMALLINT	Yes
INTEGER	INTEGER	Yes
BIGINT	BIGINT	Yes
REAL	REAL	Yes
DOUBLE PRECISION	DOUBLE PRECISION	Yes
NUMERIC	NUMERIC	Yes
CHARACTER VARYING	CHARACTER VARYING	Yes
CHARACTER	CHARACTER	Yes

Data Type (PostgreSQL)	Data Type (GaussDB)	Whether to Support Mapping
BIT	BIT	Yes
BIT VARYING	BIT VARYING	Yes
BOOLEAN	BOOLEAN	Yes
BYTEA	BYTEA	Yes
TEXT	TEXT	Yes
TIME WITHOUT TIME ZONE	TIME WITHOUT TIME ZONE	Yes
TIME WITH TIME ZONE	TIME WITH TIME ZONE	Yes
TIMESTAMP WITHOUT TIME ZONE	TIMESTAMP WITHOUT TIME ZONE	Yes
TIMESTAMP WITH TIME ZONE	TIMESTAMP WITH TIME ZONE	Yes
INTERVAL	INTERVAL	Yes
CIDR	CIDR	Yes
PATH	PATH	Yes
BOX	BOX	Yes
LSEG	LSEG	Yes
MACADDR	MACADDR	Yes
POINT	POINT	Yes
POLYGON	POLYGON	Yes
INET	INET	Yes
TSQUERY	TSQUERY	Yes
TSVECTOR	TSVECTOR	Yes
UUID	UUID	Yes
JSON	JSON	Yes
JSONB	JSONB	Yes

8.16 TiDB->TaurusDB

Table 8-19 Data type mapping

Data Type (TiDB)	Data Type (TaurusDB)	Whether to Support Mapping
BIGINT	BIGINT	Yes
BINARY	BINARY	Yes
BIT	BIT	Yes
BLOB	BLOB	Yes
BOOLEAN	BOOLEAN	Yes
CHAR	CHAR	Yes
DATE	DATE	Yes
DATETIME	DATETIME	Yes
DECIMAL	DECIMAL	Yes
DOUBLE	DOUBLE	Yes
ENUM	ENUM	Yes
FLOAT	FLOAT	Yes
INT	INT	Yes
JSON	JSON	Yes
LOB	LOB	Yes
LONGTEXT	LONGTEXT	Yes
MEDIUMBLOB	MEDIUMBLOB	Yes
MEDIUMINT	MEDIUMINT	Yes
SET	SET	Yes
SMALLINT	SMALLINT	Yes
TEXT	TEXT	Yes
TIME	TIME	Yes
TIMESTAMP	TIMESTAMP	Yes
TINYBLOB	TINYBLOB	Yes
TINYINT	TINYINT	Yes
TINYTEXT	TINYTEXT	Yes

Data Type (TiDB)	Data Type (TaurusDB)	Whether to Support Mapping
VARBINARY	VARBINARY	Yes
VARCHAR	VARCHAR	Yes
YEAR	YEAR	Yes

8.17 Microsoft SQL Server->GaussDB

Table 8-20 Data type mapping

Data Type (Microsoft SQL Server)	Data Type (GaussDB)	Whether to Support Mapping
TINYINT	SMALLINT	Yes
SMALLINT	SMALLINT	Yes
INT	INTEGER	Yes
BIGINT	BIGINT	Yes
DECIMAL	NUMERIC	Yes
NUMERIC	NUMERIC	Yes
FLOAT	DOUBLE PRECISION	Yes
REAL	REAL	Yes. The precision of SQL Server is 7 digits, while that of GaussDB is 6 digits. If the REAL value of the source database has 7 digits, 1-digit precision loss will occur when the value is synchronized to the destination database.
SMALLMONEY	NUMERIC(10,4)	Yes
MONEY	NUMERIC(19,4)	Yes
BIT	BOOLEAN	Yes
DATE	DATE	Yes
SMALLDATETIME	SMALLDATETIME	Yes
DATETIME	TIMESTAMP(3) WITHOUT TIME ZONE	Yes

Data Type (Microsoft SQL Server)	Data Type (GaussDB)	Whether to Support Mapping
DATETIME2	TIMESTAMP WITHOUT TIME ZONE	Yes
DATETIMEOFFSET	TIMESTAMP WITH TIME ZONE	Yes
TIME(p)	TIME(P) WITHOUT TIME ZONE	The value can be accurate to seconds. Decimals are discarded.
TIMESTAMP	BYTEA	Yes
XML	CLOB	Yes
CHAR	CHARACTER	Yes. If a column of this type in the source database contains characters that occupy more than one byte, increase the length of the column in the destination database.
VARCHAR	CHARACTER VARYING()	Supported
VARCHAR(max)	CLOB	Supported
NCHAR	CHARACTER VARYING()	Supported
NVARCHAR	NVARCHAR2	Yes. If a column of this type in the source database contains characters that occupy more than one byte, increase the length of the column in the destination database.
NVARCHAR(max)	CLOB	Yes
BINARY	BYTEA	Yes
VARBINARY	BYTEA	Yes
VARBINARY(max)	BLOB	Yes
IMAGE	BLOB	Yes
HIERARCHYID	BYTEA	Yes
NTEXT	CLOB	Yes
TEXT	CLOB	Yes
UNIQUEIDENTIFIER	CHARACTER(36)	Yes

8.18 Microsoft SQL Server->MySQL

Data Type (Microsoft SQL Server)	Data Type (MySQL)	Whether to Support Mapping
TINYINT	TINYINT	Yes. RDS for MySQL uses unsigned data.
SMALLINT	SMALLINT	Yes
INT	INT	Yes
BIGINT	BIGINT	Yes
DECIMAL	DECIMAL	Yes
NUMERIC	DECIMAL	Yes
FLOAT	FLOAT	Partially supported. The precision may be lost. You are not advised to use this data type as the primary key. Otherwise, there may be data inconsistency.
REAL	DOUBLE	Yes
SMALLMONEY	DECIMAL(10,4)	Yes. The currency symbol will be lost.
MONEY	DECIMAL(19,4)	Yes. The currency symbol will be lost.
BIT	BIT	Yes
DATE	DATE	Yes. You are not advised to use this data type as the primary key. Otherwise, there may be data inconsistency.
SMALLDATETIME	DATETIME	Yes
DATETIME2	DATETIME	Yes. The precision may be lost.
DATETIME	DATETIME	Yes
DATETIMEOFFSET	TIMESTAMP	Partially supported. The time zone will be lost.
TIME	TIME	Yes
XML	LONGTEXT	Yes

Data Type (Microsoft SQL Server)	Data Type (MySQL)	Whether to Support Mapping
CHAR	CHAR/VARCHAR	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is less than 256, the converted data type is CHAR. If the defined character length is greater than or equal to 256, the converted data type is VARCHAR.
VARCHAR	VARCHAR/LONGTEXT	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARCHAR. If the defined character length is max, the converted data type is LONGTEXT.
BINARY	BINARY/BLOB	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is less than 256, the converted data type is BINARY. If the defined character length is greater than or equal to 256, the converted data type is BLOB.

Data Type (Microsoft SQL Server)	Data Type (MySQL)	Whether to Support Mapping
VARBINARY	VARBINARY/LONG BLOB	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARBINARY. If the defined character length is max, the converted data type is LONGBLOB.
IMAGE	LONGBLOB	Yes
NTEXT	LONGTEXT	Yes
TEXT	LONGTEXT	Yes
NCHAR	CHAR/VARCHAR	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is less than 256, the converted data type is CHAR. If the defined character length is greater than or equal to 256, the converted data type is VARCHAR.
NVARCHAR	VARCHAR/ LONGTEXT	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARCHAR. If the defined character length is max, the converted data type is LONGTEXT.
UNIQUEIDENTIFIER	CHAR(36)	Yes
TIMESTAMP	BINARY(8)	Yes

 NOTE

- After **datetimeoffset** of Microsoft SQL Server is converted to **timestamp** of MySQL, the time zone is lost.

The value of **datetimeoffset** of Microsoft SQL Server ranges from 0001-01-01 00:00:00.0000000 to 9999-12-31 23:59:59.9999999 (UTC time). The value of **timestamp** of MySQL ranges from 1970-01-01 00:00:01.000000 to 2038-01-19 03:14:07.999999. The date type range of Microsoft SQL Server is larger than that of MySQL. If a value is out of the range, DRS reports an error by default.

- After **datetime2** of Microsoft SQL Server is converted to **datetime** of MySQL, the maximum value decreases.

The maximum value of Microsoft SQL Server is 9999-12-31 23:59:59.999999, and that of MySQL is 9999-12-31 23:59:59.499999. DRS processes the maximum value as 9999-12-31 23:59:59.

8.19 Microsoft SQL Server->TaurusDB

Data Type (Microsoft SQL Server)	Data Type (TaurusDB)	Whether to Support Mapping
TINYINT	TINYINT	Yes. TaurusDB uses unsigned data.
SMALLINT	SMALLINT	Yes
INT	INT	Yes
BIGINT	BIGINT	Yes
DECIMAL	DECIMAL	Yes
NUMERIC	DECIMAL	Yes
FLOAT	FLOAT	Partially supported. The precision may be lost. You are not advised to use this data type as the primary key. Otherwise, there may be data inconsistency.
REAL	DOUBLE	Yes
SMALLMONEY	DECIMAL(10,4)	Yes. The currency symbol will be lost.
MONEY	DECIMAL(19,4)	Yes. The currency symbol will be lost.
BIT	BIT	Yes
DATE	DATE	Yes. You are not advised to use this data type as the primary key. Otherwise, there may be data inconsistency.

Data Type (Microsoft SQL Server)	Data Type (TaurusDB)	Whether to Support Mapping
SMALLDATETIME	DATETIME	Yes
DATETIME2	DATETIME	Yes. The precision may be lost.
DATETIME	DATETIME(3)	Yes
DATETIMEOFFSET	TIMESTAMP	Partially supported. The time zone will be lost.
TIME	TIME	Yes
XML	LONGTEXT	Yes
CHAR	CHAR/VARCHAR	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is less than 256, the converted data type is CHAR. If the defined character length is greater than or equal to 256, the converted data type is VARCHAR.
VARCHAR	VARCHAR/LONGTEXT	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARCHAR. If the defined character length is max, the converted data type is LONGTEXT.
BINARY	BINARY/BLOB	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is less than 256, the converted data type is BINARY. If the defined character length is greater than or equal to 256, the converted data type is BLOB.

Data Type (Microsoft SQL Server)	Data Type (TaurusDB)	Whether to Support Mapping
VARBINARY	VARBINARY/LONG BLOB	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARBINARY. If the defined character length is max, the converted data type is LONGBLOB.
IMAGE	LONGBLOB	Yes
NTEXT	LONGTEXT	Yes
TEXT	LONGTEXT	Yes
NCHAR	CHAR/VARCHAR	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is less than 256, the converted data type is CHAR. If the defined character length is greater than or equal to 256, the converted data type is VARCHAR.
NVARCHAR	VARCHAR/ LONGTEXT	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARCHAR. If the defined character length is max, the converted data type is LONGTEXT.
UNIQUEIDENTIFIER	CHAR(36)	Yes
TIMESTAMP	BINARY(8)	Yes

 NOTE

- After **datetimeoffset** of Microsoft SQL Server is converted to **timestamp** of TaurusDB, the time zone is lost.

The value of **datetimeoffset** of Microsoft SQL Server ranges from 0001-01-01 00:00:00.0000000 to 9999-12-31 23:59:59.9999999 (UTC time). The value of **timestamp** of TaurusDB ranges from 1970-01-01 00:00:01.000000 to 2038-01-19 03:14:07.999999. The date type range of Microsoft SQL Server is larger than that of TaurusDB. If a value is out of the range, DRS reports an error by default.

- After **datetime2** of Microsoft SQL Server is converted to **datetime** of TaurusDB, the maximum value decreases.

The maximum value of Microsoft SQL Server is 9999-12-31 23:59:59.999999, and that of MySQL is 9999-12-31 23:59:59.499999. DRS processes the maximum value as 9999-12-31 23:59:59.

8.20 Microsoft SQL Server->PostgreSQL

Data Type (Microsoft SQL Server)	Data Type (PostgreSQL)	Whether to Support Mapping
TINYINT	INT2	Yes
SMALLINT	INT2	Yes
INT	INT4	Yes
BIGINT	INT8	Yes
DECIMAL	NUMERIC	Yes
NUMERIC	NUMERIC	Yes
FLOAT	FLOAT8	Yes
REAL	FLOAT4	Yes. The precision of SQL Server is 7 digits, while that of PostgreSQL is 6 digits. If the REAL value of the source database has 7 digits, 1-digit precision loss will occur when the value is synchronized to the destination database.
SMALLMONEY	NUMERIC (10,4)	Yes. The currency symbol will be lost.
MONEY	NUMERIC (19,4)	Yes. The currency symbol will be lost.
BIT	BOOL	Yes
DATE	DATE	Yes
SMALLDATETIME	TIME(0)	Yes

Data Type (Microsoft SQL Server)	Data Type (PostgreSQL)	Whether to Support Mapping
DATETIME2	TIME(6)	Yes. The precision may be lost.
DATETIME	TIME (3)	Yes
DATETIMEOFFSET	TIMESTAMPTZ(6)	Partially supported.
TIME(p)	TIME(p)	Yes. The maximum precision of PostgreSQL is 6.
XML	TEXT	Yes
CHAR[(M)]	CHAR	Yes
VARCHAR(M)	VARCHAR/TEXT	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARCHAR. If the defined character length is max, the converted data type is TEXT. SQL Server supports a maximum of 4 GB, and PostgreSQL supports only 2 GB.
BINARY[(M)]	BYTEA	Yes
VARBINARY(M)	BYTEA	Yes
IMAGE	BYTEA	Yes
NTEXT	TEXT	Yes. SQL Server supports a maximum of 4 GB, and PostgreSQL supports only 2 GB.
TEXT	TEXT	Yes. SQL Server supports a maximum of 4 GB, and PostgreSQL supports only 2 GB.
NCHAR	CHAR	Yes

Data Type (Microsoft SQL Server)	Data Type (PostgreSQL)	Whether to Support Mapping
NVARCHAR(M)	VARCHAR/TEXT	Yes. Data is converted into the corresponding type based on the defined character length. If the defined character length is a specific number, the converted data type is VARCHAR. If the defined character length is max, the converted data type is TEXT. SQL Server supports a maximum of 4 GB, and PostgreSQL supports only 2 GB.
UNIQUEIDENTIFIER	VARCHAR(36)	Yes
TIMESTAMP	BYTEA (8)	Yes

9 Security

9.1 Shared Responsibilities

Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Unlike traditional on-premises data centers, cloud computing separates operators from users. This approach not only enhances flexibility and control for users but also greatly reduces their operational workload. For this reason, cloud security cannot be fully ensured by one party. Cloud security requires joint efforts of Huawei Cloud and you, as shown in [Figure 9-1](#).

- **Huawei Cloud:** Huawei Cloud is responsible for infrastructure security, including security and compliance, regardless of cloud service categories. The infrastructure consists of physical data centers, which house compute, storage, and network resources, virtualization platforms, and cloud services Huawei Cloud provides for you. In PaaS and SaaS scenarios, Huawei Cloud is responsible for security settings, vulnerability remediation, security controls, and detecting any intrusions into the network where your services or Huawei Cloud components are deployed.
- **Customers:** As our customer, your ownership of and control over your data assets will not be transferred under any cloud service category. Without your explicit authorization, Huawei Cloud will not use or monetize your data, but you are responsible for protecting your data and managing identities and access. This includes ensuring the legal compliance of your data on the cloud, using secure credentials (such as strong passwords and multi-factor authentication), and properly managing those credentials, as well as monitoring and managing content security, looking out for abnormal account behavior, and responding to it, when discovered, in a timely manner.

Figure 9-1 Huawei Cloud shared security responsibility model

Cloud security responsibilities are determined by control, visibility, and availability. When you migrate services to the cloud, assets, such as devices, hardware, software, media, VMs, OSs, and data, are controlled by both you and Huawei Cloud. This means that your responsibilities depend on the cloud services you select. As shown in [Figure 9-1](#), customers can select different cloud service types (such as IaaS, PaaS, and SaaS services) based on their service requirements. As control over components varies across different cloud service categories, the responsibilities are shared differently.

- In on-premises scenarios, customers have full control over assets such as hardware, software, and data, so tenants are responsible for the security of all components.
- In IaaS scenarios, customers have control over all components except the underlying infrastructure. So, customers are responsible for securing these components. This includes ensuring the legal compliance of the applications, maintaining development and design security, and managing vulnerability remediation, configuration security, and security controls for related components such as middleware, databases, and operating systems.
- In PaaS scenarios, customers are responsible for the applications they deploy, as well as the security settings and policies of the middleware, database, and network access under their control.
- In SaaS scenarios, customers have control over their content, accounts, and permissions. They need to protect their content, and properly configure and protect their accounts and permissions in compliance with laws and regulations.

On-premises (On-Prem): Software and IT infrastructure that are deployed and managed by customers within their own data centers, rather than be deployed by remote cloud service providers.

Infrastructure as a Service (IaaS): Cloud service providers offer compute, network, storage, and more infrastructure services, including [Elastic Cloud Server \(ECS\)](#), [Virtual Private Network \(VPN\)](#), and [Object Storage Service \(OBS\)](#).

Platform as a Service (PaaS): Cloud service providers deliver platforms required for application development and deployment, such as [ModelArts](#) and [GaussDB](#). Customers do not need to maintain the underlying infrastructure.

Software as a Service (SaaS): Cloud service providers offer complete application software, such as [Huawei Cloud Meeting](#). Customers use the software directly without the need to install the application, maintain it, or manage its underlying platform or infrastructure.

9.2 Identity Authentication and Access Control

Identity Authentication

DRS uses [Identity and Access Management \(IAM\)](#) to implement fine-grained permission management. IAM provides identity authentication, permission assignment and access control, grants different permissions to different user groups, uses fine-grained authentication to control the usage scope of DRS resources, and ensures users have secure access to resources. For details, see [Creating an IAM User and Logging In](#).

Access Control

- **Permissions control**

You can use IAM to assign different permissions to different employees in your enterprise to access your instance resources. For details about DRS permissions, see [Permissions Management](#).

- **Network isolation**

When creating a DRS instance, you can select a subnet in the VPC where the DRS instance is located based on service requirements. After the DRS instance is created, DRS will assign an IP address in the subnet to the DRS instance for connecting to source and destination instances. If the DRS instance is in the same VPC as the source instance or destination instance on Huawei Cloud, you can configure security groups for the source instance, destination instance, or DRS instance to control network access.

For details, see [Creating a VPC](#).

9.3 Data Protection

DRS provides a series of methods and features to ensure data security and integrity during transmission.

Host Security and Data Reliability and Durability

At the underlying layer, DRS uses [Elastic Cloud Servers \(ECSs\)](#) for computing and [Elastic Volume Service \(EVS\)](#) disks for storage. With secure ECSs and reliable EVS disks, the host security, data reliability, and data durability of DRS instances can be effectively ensured.

Instance High Availability

To improve service availability and resilience, DRS provides resumable data transfer and fault recovery. If data in the source database is not corrupted or lost,

the DRS instance can resume data transfer from the point at which the transfer was stopped. If the underlying resources of an instance are faulty, data is migrated to a new instance in the AZ, and then the interrupted transfer continues. DRS also provides the cross-AZ HA. If the instance in the primary AZ becomes faulty, services can be switched over to the instance in the standby AZ to continue data replication.

Data Transmission Encryption

To secure data replication, DRS allows you to encrypt data transmission over a public network, VPN, Direct Connect, or VPC.

Permanent Data Deletion

When a DRS instance is deleted, the computing and storage resources of the instance are reclaimed. In addition, all data on the DRS instance is deleted and cannot be restored, including basic instance information, run logs, and data comparison results.

9.4 Audit and Logs

Audit

Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, track resource changes, audit compliance, and locate faults.

For details about how to enable and configure CTS, see [CTS Overview](#).

With CTS, you can record operations associated with DRS for later query, audit, and backtracking. For details, see [Key Operations Recorded by CTS](#).

Logs

DRS logs refer to the warning-, error-, and info-level logs generated during the migration process. You can view logs to locate and analyze database problems and rectify tasks. For details, see [Viewing Migration Logs](#).

9.5 Risk Monitoring

Cloud Eye is a comprehensive monitoring platform for resources like cloud databases and cloud servers. It enables you to monitor resources, configure alarm rules, identify resource exceptions, and quickly respond to resource changes.

You can monitor resources and operations, such as CPU usage and network throughput using Cloud Eye. You can configure DRS alarm rules to customize the monitored objects and notification policies and learn the DRS running status in a timely manner. For details about supported monitoring metrics and how to create alarm rules, see [Supported Metrics](#).

Cloud Eye can monitor performance metrics from the last 1 hour, 3 hours, 12 hours, 1 day, 7 days, and 6 months.

9.6 Fault Recovery

Resumable Data Transfer

To improve service availability and resilience, DRS provides resumable data transfer and fault recovery. If data in the source database is not corrupted or lost, the DRS instance can resume data transfer from the point at which the transfer was stopped.

Instance High Availability

An availability zone (AZ) is a physical region where resources use independent power supply and networks. AZs are physically isolated but interconnected over a local network. DRS also provides the cross-AZ HA. If the instance in the primary AZ becomes faulty, services can be switched over to the instance in the standby AZ to continue data replication.

9.7 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO), system and organization controls (SOC), and Payment card industry (PCI) compliance standards. These certifications are available for [download](#).

Figure 9-2 Downloading compliance certificates

Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see [Resource Center](#).

Figure 9-3 Resource Center

9.8 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.

10 Permissions Management

If you need to assign different permissions to employees in your enterprise to access your DRS resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you to securely access your Huawei Cloud resources.

With IAM, you can use your Huawei account to create IAM users for your employees, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use DRS but must not delete DRS resources or perform any high-risk operations. To achieve this result, you can create IAM users for the software developers and grant them only the permissions required for using DRS resources.

If your Huawei account does not need individual IAM users for permissions management, you may skip over this topic.

IAM can be used free of charge. You pay only for the resources in your Huawei Cloud account. For more information about IAM, see [IAM Service Overview](#).

DRS Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

DRS is a project-level service deployed and accessed in specific physical regions. To assign DRS permissions to a user group, specify the scope as region-specific (for example, CN-Hong Kong) projects and select projects (for example, ap-southeast-1) for the permissions to take effect. If you select **All projects**, the permissions will take effect for user groups in all region-specific projects. When accessing DRS, users need to switch to a region where they have been authorized to use DRS.

You can grant users permissions by using roles and policies.

- **Roles:** A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you need to also assign other roles which the permissions depend on to take effect. However, roles are not an ideal choice for fine-grained authorization and secure access control.

- **Policies:** A type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, you can grant ECS users only the permissions for managing a certain type of ECSs. Most fine-grained policies are API-based.

Table 10-1 lists all the system policies supported by DRS.

Table 10-1 System-defined roles and policies supported by DRS

Policy Name/ System Role	Description	Type	Dependency
Security Administrator	Security administrator 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. If the automatic function is unavailable, see Why Cannot Scheduled DRS Tasks Be Started?	System-defined role	None
DRS Administrator	DRS administrator Basic permission, which must be added when DRS is used.	System-defined role	Dependent on the Tenant Guest, Server Administrator, and RDS Administrator roles. <ul style="list-style-type: none"> • Tenant Guest: A project-level role, which must be assigned in the same project. • Server Administrator: A project-level role, which must be assigned in the same project. • RDS Administrator: A project-level role, which must be assigned in the same project.

Policy Name/ System Role	Description	Type	Dependency
DRS FullAccess	Full permissions for DRS	System policy	<p>Dependent on the VPC FullAccess, RDS ReadOnlyAccess, and SMN Administrator, OBS Administrator, and EPS ReadOnlyAccess policies.</p> <ul style="list-style-type: none"> • VPC FullAccess: This parameter needs to be configured when the VPC and subnet are selected. • RDS ReadOnlyAccess: This parameter needs to be configured when RDS is selected. • SMN Administrator: This parameter needs to be configured when SMN is selected. • OBS Administrator: This parameter needs to be configured when bucket information is selected for a backup task. • EPS ReadOnlyAccess: This parameter needs to be configured when an enterprise project is selected. <p>For a yearly/monthly task, the following permissions must be configured: BSS Operator or BSS Administrator</p>

Policy Name/ System Role	Description	Type	Dependency
DRS ReadOnlyAccess	Read-only permissions for DRS resources.	System policy	Configure the following policies as required: RDS ReadOnlyAccess: This parameter needs to be configured when RDS is selected. SMN Administrator: This parameter needs to be configured when SMN is selected.
DRS FullWithoutDeletePermission	All permissions on DRS except the deletion permission	System Policy	Dependent on the VPC FullAccess , RDS ReadOnlyAccess , and SMN Administrator , and OBS Administrator policies. <ul style="list-style-type: none"> • VPC FullAccess: This parameter needs to be configured when the VPC and subnet are selected. • RDS ReadOnlyAccess: This parameter needs to be configured when RDS is selected. • SMN Administrator: This parameter needs to be configured when SMN is selected. • OBS Administrator: This parameter needs to be configured when bucket information is selected for a backup task. For a yearly/monthly task, the following permissions must be configured: BSS Operator or BSS Administrator

 **NOTE**

In addition to the preceding permissions, the read permission for the corresponding DB instance is required. For example, if a DDM database is used, configure the DDM ReadOnlyAccess permission for the project. If a DDS database is used, configure the DDS ReadOnlyAccess permission for the project.

Table 10-2 lists the common operations supported by the DRS system policy.

Table 10-2 Common operations supported by the DRS system policy

Procedure	DRS FullAccess	DRS ReadOnlyAccess	DRS Administrator	DRS FullWithOutDeletePermission
Creating a task	√	x	√	√
Editing a task	√	x	√	√
Deleting a task	√	x	√	x
Starting a task	√	x	√	√
Retrying a task	√	x	√	√
Stopping a task	√	x	√	√

Table 10-3 lists common DRS operations and corresponding actions. You can refer to this table to customize permission policies.

Table 10-3 Common operations and supported actions

Permission	Actions	Remarks
Querying the RPO and RTO	drs:dataGuardJob:list	None
Performing a primary/standby switchover	drs:disasterRecovery-Job:switchover	None

Permission	Actions	Remarks
Performing operations on tasks.	drs:migrationJob:action	<p>The VPC FullAccess permission for the project is required.</p> <p>The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project. If the DDS database is used, configure the DDS ReadOnlyAccess permission for the project.</p> <p>The OBS Administrator permission for backup tasks is required.</p>
Stopping a task	drs:migrationJob:terminate	<p>Permissions required for the project:</p> <p>VPC FullAccess</p> <p>RDS ReadOnlyAccess</p> <p>Permissions required for the backup task:</p> <p>OBS Administrator</p> <p>Permission required for subscribing to message notification:</p> <p>SMN Administrator</p>
Modifying a migration task	drs:migrationJob:modify	<p>Permission required for selecting VPCs and subnets on the GUI:</p> <p>VPC FullAccess</p> <p>The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project. If the DDS database is used, configure the DDS ReadOnlyAccess permission for the project.</p> <p>Permission required for subscribing to message notification:</p> <p>SMN Administrator</p>

Permission	Actions	Remarks
<p>Creating a migration task</p> <p>Cloning a task</p>	drs:migrationJob:create	<p>Permission required for selecting VPCs and subnets on the GUI:</p> <p>VPC FullAccess</p> <p>The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project. If the DDS database is used, configure the DDS ReadOnlyAccess permission for the project.</p> <p>Permission required for subscribing to message notification:</p> <p>SMN Administrator</p> <p>For a yearly/monthly task, the following permissions must be configured:</p> <p>BSS Operator or BSS Administrator</p>
Deleting a migration task	drs:migrationJob:delete	None
Updating the database user information.	drs:migrationJob:modifyUserInfo	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Controlling the migration speed	drs:migrationJob:setMigrationTransSpeed	None
Modify database parameters	drs:databaseParams:modify	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.

Permission	Actions	Remarks
Updating the data processing information	drs:dataTransformation:update	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Adding the data processing information	drs:dataTransformation:add	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Deleting the data processing data	drs:dataTransformation:delete	None
Updating the database object selection information	drs:migrationJob:update	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Updating the ignore policy of a synchronization task	drs:synchronizationJob:update	None
Updating the task configuration	drs:migrationJob:updateJobConfig	None
Updating the DDL filtering policy.	drs:migrationJob:updateDDLPolicy	None
Modifying the comparison policy	drs:healthCompare:modify	None
Stopping a comparison task	drs:healthCompare:stop	None
Creating an object-level table comparison task	drs:migrationCompareJob:create	None
Canceling a data-level table comparison task	drs:migrationCompareJob:delete	None
Immediately starting a data-level table comparison task	drs:migrationCompareJob:start	None

Permission	Actions	Remarks
Creating a data subscription task	drs:subscriptionJob:create	Permission required for selecting RDS on the GUI: RDS ReadOnlyAccess Permission required for subscribing to message notification: SMN Administrator Permission required for the project: VPC FullAccess
Modifying a data subscription task	drs:subscriptionJob:update	Permission required for selecting RDS on the GUI: OBS Administrator RDS ReadOnlyAccess Permission required for subscribing to message notification: SMN Administrator Permission required for the project: VPC FullAccess
Editing the data subscription information	drs:subscriptionJob:edit	Permission required for selecting buckets and RDS on the GUI: OBS Administrator RDS ReadOnlyAccess Permission required for subscribing to message notification: SMN Administrator Permissions required: VPC FullAccess
Deleting a subscription task	drs:subscriptionJob:delete	None

Permission	Actions	Remarks
Performing operations on a subscription task	drs:subscriptionJob:subscribe	Permissions required for the project: VPC FullAccess RDS ReadOnlyAccess Permissions required for the backup task: OBS Administrator Permission required for subscribing to message notification: SMN Administrator
Modifying consumption start time	drs:subscriptionJob:UpdateConsumeTime	None
Creating a backup migration task 2.0	drs:backupMigration-Job:create	Permission required for selecting buckets and RDS on the GUI: OBS Administrator RDS ReadOnlyAccess Permission required for subscribing to message notification: SMN Administrator
Deleting a backup migration task	drs:backupMigration-Job:delete	None
Modifying backup migration task details	drs:backupMigration-Job:modify	Permission required for selecting buckets and RDS on the GUI: OBS Administrator RDS ReadOnlyAccess Permission required for subscribing to message notification: SMN Administrator
Freezing and unfreezing resources.	drs:frozenOrUnfreeze-Job:frozen	None
Cleaning up resources	drs:cleanJob:clean	The VPC FullAccess permission is required.
Verifying the backup task name.	drs:backupMigration-Job:check	None

Permission	Actions	Remarks
Verifying data processing	drs:dataTransformation:check	None
Verifying online task names	drs:migrationJob:check	None
Obtaining database parameters	drs:databaseParameters:get	None
Querying backup migration tasks	drs:backupMigrationJob:list	None
Querying backup migration task details	drs:backupMigrationJob:get	None
Obtaining the data subscription task details	drs:subscriptionJob:get	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Querying operation results	drs:job:getResult	None
Obtaining the task pricing details	drs:migrationJob:getMeteringPrice	None
Querying the data processing information	drs:migrationTransformationJob:get	None
Obtaining the task pre-check results	drs:precheckJob:get	None
Obtaining the object-level migration comparison overview	drs:compareJob:getOverview	None
Querying data-level table comparison tasks	drs:compareJob:list	None
Querying data-level table comparison results	drs:compareJob:getResult	None
Obtaining object-level migration comparison details	drs:compareJob:getDetails	None

Permission	Actions	Remarks
Querying details about a data-level table comparison task	drs:compareJob:getContentsInfo	None
Querying the estimated time of a comparison task	drs:compareJob:getEstimateTime	None
Querying the value comparison overview.	drs:compareJob:getContentOverview	None
Querying the row comparison overview	drs:compareJob:getLineOverview	None
Querying row comparison details	drs:compareJob:getLineDetail	None
Obtaining account comparison details	drs:compareJob:getAccountDetails	None
Querying value comparison details	drs:compareJob:getContentDetail	None
Querying value comparison differences	drs:compareJob:getContentDiff	None
Obtaining the online migration task list	drs:migrationJob:list	None
Obtaining the online migration task details	drs:migrationJob:get	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Obtaining the object-level migration comparison overview	drs:migrationJob:getCompareStruct	None
Obtaining the data-level stream comparison	drs:migrationJob:getStreamComparison	None

Permission	Actions	Remarks
Obtaining the source database user list	drs:migrationJob:getSrcUsers	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Obtaining the migration progress of a specified migration task	drs:migrationJob:getSpecifiedProgress	None
Obtaining the database affected time of a specified task.	drs:migrationJob:getEffectTime	None
Querying the migration progress	drs:migrationJobs:getProgress	None
Querying the health comparison report list	drs:healthCompareJob:list	None
Obtaining the object-level migration comparison overview	drs:healthCompareJob:getOverview	None
Obtaining object-level comparison details	drs:healthCompareJob:getObjectDetail	None
Obtaining account comparison details	drs:healthCompareJob:getAccountDetails	None
Querying row comparison details	drs:healthCompareJob:getLineDetail	None
Querying the comparison policy	drs:healthCompareJob:getComparePolicy	None
Obtaining the disaster recovery monitoring data	drs:disasterRecoveryJob:get	Permissions required: CES ReadOnlyAccess
Obtaining the RPO and RTO of a specified task	drs:dataGuardJob:list	None
Obtaining the replay failure list	drs:replayFaultsJob:list	None

Permission	Actions	Remarks
Processing data	drs:migrationJob:action	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Starting a task	drs:migrationJob:action	The VPC FullAccess permission is required.
Querying task details	drs:migrationJob:get	The read permission for the corresponding instance is required. For example, if the RDS database is used, configure the RDS ReadOnlyAccess permission for the project.
Querying task statuses	drs:migrationJob:get	None
Querying resource tags	drs:tag:get	None
Deleting resource tags	drs:tags:delete	None
Adding a resource tag	drs:tag:add	None
Modifying resource tags	drs:tag:modify	None
Obtaining migration logs	drs:migrationJob:getLog	None
Querying the involved Kafka topic information	drs:kafkaJob:get	None
Obtaining the supported feature list	drs:supportFeature:get	None
Querying the feature whitelist	drs:featureWhiteJob:list	None
Querying the quotas that need adjustment	drs:quota:adjust	None
Updating the quotas	drs:quota:update	None

Permission	Actions	Remarks
Querying resource quotas	drs:quota:get	None
Updating the user guide	drs:userGuide:update	None
Obtaining user guide details	drs:userGuide:list	None
Querying predefined tags	-	To query predefined tags, configure the following action: tms:resourceTags:list
Querying configured log groups	-	To query configured log groups, configure the following action: lts:groups:get
Querying configured log streams	-	To query configured log streams, configure the following action: lts:topics:get
Obtaining the real-time synchronization task list	drs:synchronizationJob:list	None
Obtaining the real-time synchronization task details	drs:synchronizationJob:get	None
Obtaining the real-time DR task details	drs:dataGuardJob:get	None
Obtaining the data subscription task list	drs:subscriptionJob:list	None
Obtaining the workload replay task list	drs:replayJob:list	None
Obtaining the workload replay task details	drs:replayJob:get	None

11 IAM-based Permissions Management

If you need to assign different permissions to employees in your enterprise to access your DRS resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you to securely access your Huawei Cloud resources. If your Huawei Cloud account does not require IAM for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources in your account.

With IAM, you can control access to specific Huawei Cloud resources. For example, if you want some software developers in your enterprise to be able to use DRS resources but do not want them to be able to delete DRS resources or perform any other high-risk operations, you can create IAM users and grant permission to use DRS resources but not permission to delete them.

IAM supports role/policy-based authorization and identity policy-based authorization.

The following table describes the differences between these two authorization models.

Table 11-1 Differences between role/policy-based and identity policy-based authorization

Authorization Model	Core Relationship	Permissions	Authorization Method	Scenario
Role/Policy	User-permission-authorization scope	<ul style="list-style-type: none"> System-defined roles System-defined policies Custom policies 	Assigning roles or policies to principals	To authorize a user, you need to add it to a user group first and then specify the scope of authorization. It provides a limited number of condition keys and cannot meet the requirements of fine-grained permissions control. This method is suitable for small- and medium-sized enterprises.
Identity policy	User-policy	<ul style="list-style-type: none"> System-defined policies Custom identity policies 	<ul style="list-style-type: none"> Assigning identity policies to principals Attaching identity policies to principals 	You can authorize a user by attaching an identity policy to it. User-specific authorization and a variety of key conditions allow for more fine-grained permissions control. However, this model can be hard to set up. It requires a certain amount of expertise and is suitable for medium- and large-sized enterprises.

Assume that you want to grant IAM users the permissions needed to create DRS tasks in CN North-Beijing4 and OBS buckets in CN South-Guangzhou. With role/policy-based authorization, the administrator needs to create two custom policies and assign both to the IAM users. With identity policy-based authorization, the administrator only needs to create one custom identity policy and configure the condition key **g:RequestedRegion** for the policy, and then attaches the policy to the users or grants the users access permissions to the specified regions. Identity policy-based authorization is more flexible than role/policy-based authorization.

Policies/identity policies and actions in the two authorization models are not interoperable. You are advised to use the identity policy-based authorization model. For details about system-defined permissions, see [Role/Policy-based Authorization](#) and [Identity Policy-based Authorization](#).

For more information about IAM, see [IAM Service Overview](#).

Role/Policy-based Authorization

DRS supports role/policy-based authorization. New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and then attach policies or roles to these groups. The users then inherit permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

DRS is a project-level service deployed in specific physical regions. When you set **Scope** to **Region-specific projects** and select the specified projects (for example, **ap-southeast-2**) in the specified regions (for example, **AP-Bangkok**), the users only have permissions for resources in the selected projects. If you set **Scope** to **All resources**, the users have permissions for resources in all region-specific projects. When accessing DRS, the users need to switch to the authorized region.

Table 11-2 lists all the system-defined permissions for DRS. System-defined policies in role/policy-based authorization are not interoperable with those in identity policy-based authorization.

Table 11-2 System-defined roles and policies supported by DRS

Policy Name/ System Role	Description	Type	Dependency
Security Administrator	Security administrator 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. If the automatic function is unavailable, see Why Cannot Scheduled DRS Tasks Be Started?	System-defined role	None

Policy Name/ System Role	Description	Type	Dependency
DRS Administrator	DRS administrator Basic permission, which must be added when DRS is used.	System-defined role	Dependent on the Tenant Guest, Server Administrator, and RDS Administrator roles. <ul style="list-style-type: none"> • Tenant Guest: A project-level role, which must be assigned in the same project. • Server Administrator: A project-level role, which must be assigned in the same project. • RDS Administrator: A project-level role, which must be assigned in the same project.

Policy Name/ System Role	Description	Type	Dependency
DRS FullAccess	Full permissions for DRS	System policy	<p>Dependent on the VPC FullAccess, RDS ReadOnlyAccess, and SMN Administrator, OBS Administrator, and EPS ReadOnlyAccess policies.</p> <ul style="list-style-type: none"> • VPC FullAccess: This parameter needs to be configured when the VPC and subnet are selected. • RDS ReadOnlyAccess: This parameter needs to be configured when RDS is selected. • SMN Administrator: This parameter needs to be configured when SMN is selected. • OBS Administrator: This parameter needs to be configured when bucket information is selected for a backup task. • EPS ReadOnlyAccess: This parameter needs to be configured when an enterprise project is selected. <p>For a yearly/monthly task, the following permissions must be configured: BSS Operator or BSS Administrator</p>

Policy Name/ System Role	Description	Type	Dependency
DRS ReadOnlyAccess	Read-only permissions for DRS resources.	System policy	Configure the following policies as required: RDS ReadOnlyAccess: This parameter needs to be configured when RDS is selected. SMN Administrator: This parameter needs to be configured when SMN is selected.
DRS FullWithoutDeletePermission	All permissions on DRS except the deletion permission	System Policy	Dependent on the VPC FullAccess , RDS ReadOnlyAccess , and SMN Administrator , and OBS Administrator policies. <ul style="list-style-type: none"> • VPC FullAccess: This parameter needs to be configured when the VPC and subnet are selected. • RDS ReadOnlyAccess: This parameter needs to be configured when RDS is selected. • SMN Administrator: This parameter needs to be configured when SMN is selected. • OBS Administrator: This parameter needs to be configured when bucket information is selected for a backup task. For a yearly/monthly task, the following permissions must be configured: BSS Operator or BSS Administrator

 NOTE

In addition to the preceding permissions, the read permission for the corresponding DB instance is required. For example, if a DDM database is used, configure the DDM ReadOnlyAccess permission for the project. If a DDS database is used, configure the DDS ReadOnlyAccess permission for the project.

Table 11-3 lists the common operations supported by system-defined permissions for DRS.

Table 11-3 Common operations supported by the DRS system policy

Procedure	DRS FullAccess	DRS ReadOnlyAccess	DRS Administrator	DRS FullWithOutDeletePermission
Creating a task	√	x	√	√
Editing a task	√	x	√	√
Deleting a task	√	x	√	x
Starting a task	√	x	√	√
Retrying a task	√	x	√	√
Stopping a task	√	x	√	√

Identity Policy-based Authorization

DRS supports identity policy-based authorization. **Table 11-4** lists all the system-defined identity policies for DRS. System-defined policies in identity policy-based authorization are not interoperable with those in role/policy-based authorization.

Table 11-4 System-defined identity policies for DRS

Identity Policy Name	Description	Type	Remarks
DRSAdministratorPolicy	Administrator permissions for DRS.	System-defined identity policy	-
DRSReadOnlyAccessPolicy	Read-only permissions for DRS.	System-defined identity policy	-

Identity Policy Name	Description	Type	Remarks
DRSFullAccessPolicy	Administrator permissions for DRS. Users granted these permissions can operate and use DRS.	System-defined identity policy	Configure required policies according to the resources you select. <ul style="list-style-type: none"> • VPCFullAccessPolicy: This policy needs to be configured when a VPC and subnet are selected. • SMNAdministratorPolicy: This policy needs to be configured when SMN is selected. • OBSAdministratorPolicy: This policy needs to be configured when bucket information is selected for a backup task. • ReadOnlyAccessPolicy: This policy needs to be configured when an enterprise project is selected.

Identity Policy Name	Description	Type	Remarks
DRSFullWithOutDeleteAccessPolicy	Full permissions for DRS, except for those for deleting tasks.	System-defined identity policy	Configure required policies according to the resources you select. <ul style="list-style-type: none"> • VPCFullAccessPolicy: This policy needs to be configured when a VPC and subnet are selected. • SMNAdministratorPolicy: This policy needs to be configured when SMN is selected. • OBSAdministratorPolicy: This policy needs to be configured when bucket information is selected for a backup task. • ReadOnlyAccessPolicy: This policy needs to be configured when an enterprise project is selected.

 **NOTE**

In addition to the above policies, you need to configure read-only policies for DB instances according to the resources you select. For example, if you select an RDS database, you need to configure RDSReadOnlyAccessPolicy. If you select a DDS database, you need to configure DDSReadOnlyAccessPolicy. If you select a DDM database, you need to configure the DDM ReadOnlyAccess permission because DDM does not support the policy-based authorization model.

Table 11-5 lists the common operations supported by system-defined identity policies for DRS.

Table 11-5 Common operations supported by system-defined identity policies

Operation	DRSFullAccessPolicy	DRSReadOnlyAccessPolicy	DRSAdministratorPolicy	DRSFullWithOutDeleteAccessPolicy
Creating a task	√	x	√	√

Operation	DRSFullAccessPolicy	DRSReadOnlyAccessPolicy	DRSAdministratorPolicy	DRSFullWithoutDeleteAccessPolicy
Editing a task	√	x	√	√
Deleting a task	√	x	√	x
Starting a task	√	x	√	√
Retrying a task	√	x	√	√
Stopping a task	√	x	√	√

Helpful Links

[IAM Service Overview](#)

12 Agency Management

If you use a member account to create a DRS task, automated tasks may fail, including scheduled startup, completion, resumable transfer, and scheduled comparison. To rectify the fault, you can use an agency to create a task.

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

Solution

- Method 1: Use the master account to create a task again because the master account has the Security Administrator permission by default. After the task is created using the master account, an agency is created.
- Method 2: Use the master account to add the Security Administrator permission to the user group to which the member account belongs, and create a task again. For details about how to add permissions, see [Creating a User Group and Assigning Permissions](#).
- Method 3: Use the master account to add a custom policy to the user group to which the member account belongs, and create a task again. The procedure is as follows:
 - a. Log in to Huawei Cloud using the master account and click **Console** in the upper right corner.
 - b. On the management console, hover the mouse pointer over the username in the upper right corner, and choose **Identity and Access Management** from the drop-down list.
 - c. On the IAM console, choose **Permissions > Policies/Roles** from the navigation pane, and click **Create Custom Policy** in the upper right corner.
 - d. Enter a policy name. Configure the following content in the **Policy Content** area. For details about how to create a custom policy, see [Creating a Custom Policy](#).
 - Creating an agency: **iam:agencies:createAgency**
 - Querying the agency list: **iam:agencies:listAgencies**

- Assigning permissions to an agency:
iam:permissions:grantRoleToAgency,
iam:permissions:grantRoleToAgencyOnProject and
iam:permissions:grantRoleToAgencyOnDomain
- Querying agency permissions: **iam:roles:listRoles**,
iam:permissions:listRolesForAgencyOnProject and
iam:permissions:listRolesForAgencyOnDomain
- e. After the policy is created, add the custom policy created in **d** to the user group to which the member account belongs. Select **Global services** for **Scope**. For details about how to add a custom policy, see [Creating a User Group and Assigning Permissions](#).
- Method 4: Manually add an agency. The procedure is as follows:

 **NOTE**

To use this method, ensure that the following two policies have been assigned to the master account:

- Querying the agency list: **iam:agencies:listAgencies**
 - Querying agency permissions: **iam:roles:listRoles**,
iam:permissions:listRolesForAgencyOnProject and
iam:permissions:listRolesForAgencyOnDomain
- a. Log in to Huawei Cloud using the master account and click **Console** in the upper right corner.
 - b. On the management console, hover the mouse pointer over the username in the upper right corner, and choose **Identity and Access Management** from the drop-down list.
 - c. In the navigation pane on the left, click **Agencies**.
 - d. In the upper right corner, click **Create Agency**.
 - e. Enter **DRS_AGENCY** in field **Agency Name**. If you select **Account** for **Agency Type**, enter **op_svc_rds** in field **Delegated Account**. If you select **Cloud service** for **Agency Type**, select **MySQL** for **Delegated Account**. Select **Unlimited** for **Validity Period** and then click **OK**.

Figure 12-1 Creating an agency

The screenshot shows a form for creating an agency. It contains the following elements:

- * Agency Name:** A text input field containing "DRS_AGENCY".
- * Agency Type:** Two radio button options: "Account" (selected) and "Cloud service". Below "Account" is the text "Delegate another Huawei Cloud account to perform operations on your resources." Below "Cloud service" is the text "Delegate a cloud service to access your resources in other cloud services."
- * Delegated Account:** A text input field containing "op_svc_rds".
- * Validity Period:** A dropdown menu showing "Unlimited".
- Description:** A text area with the placeholder "Enter a brief description." and a character count "0/255".
- Buttons:** "OK" and "Cancel" buttons at the bottom.

f. On the **Select Policy/Role** page, select policies based on the task type and click **Next** in the lower right corner.

- Self-built tasks: **DRS FullAccess**
- To-the-cloud or out-of-cloud tasks: **DRS FullAccess** and **ReadOnlyAccess** of the source and destination databases

For example, for a to-the-cloud migration task from MySQL to MySQL, you need to select the **DRS FullAccess** and **RDS ReadOnlyAccess** policies.


For details about the permissions required by different DB engines, see [Table 12-1](#).

Table 12-1 DB engines and policies

DB Engine	Authorization Policy
MySQL	RDS ReadOnlyAccess
Microsoft SQL Server	RDS ReadOnlyAccess
PostgreSQL	RDS ReadOnlyAccess
MongoDB	DDS ReadOnlyAccess
DDS	DDS ReadOnlyAccess
TaurusDB	GaussDBReadOnlyAccess
DDM	DDMReadOnlyAccess and RDSReadOnlyAccess

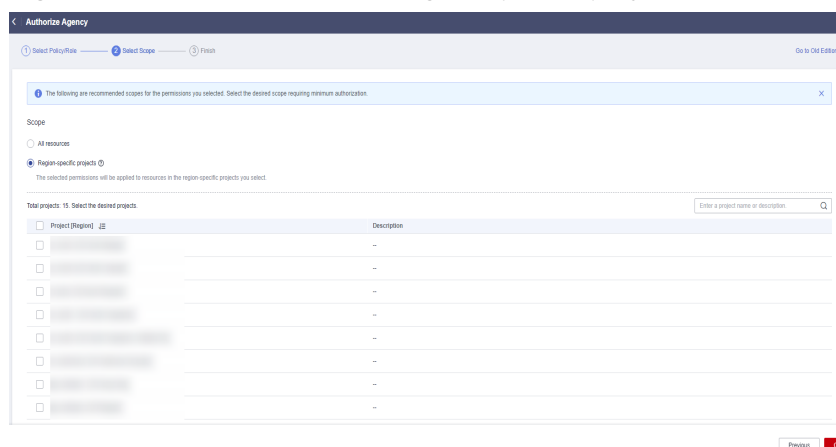
DB Engine	Authorization Policy
GaussDB Distributed	GaussDBReadOnlyAccess
GaussDB Centralized	GaussDBReadOnlyAccess
GeminiDB Mongo	GeminiDBReadOnlyAccess
Cassandra	GeminiDBReadOnlyAccess
GeminiDB Redis	GeminiDBReadOnlyAccess
MariaDB	RDSReadOnlyAccess
GeminiDB Cassandra	GeminiDBReadOnlyAccess

Figure 12-2 Select Policy/Role



- g. On the **Select Scope** page, select the scope for minimum authorization and click **OK**.
 - You need to select **Global services** and then **Region-specific projects** for the **DRS FullAccess** permission.
 - You need to select **Region-specific projects** for the **ReadOnlyAccess** permission of the source and destination databases.

Figure 12-3 Authorization for region-specific projects



- h. Click the agency name. On the **Permissions** tab, you can view permissions for the agency.

- i. The authorization takes effect after 15 to 30 minutes.

13 Notes and Constraints

13.1 Constraints on Migration Tasks

To ensure that data migration tasks run properly, DRS has certain constraints on different data flow scenarios.

MySQL Serving as the Source

To the Cloud

- [From MySQL to MySQL](#)
- [From MySQL to DDM](#)
- [From MySQL to TaurusDB](#)

From the Cloud

- [From MySQL to MySQL](#)

MySQL Schema and Logic Table Serving as the Source

- [From MySQL Schema and Logic Table to DDM](#)

MongoDB Serving as the Source

- [From MongoDB to DDS](#)
- [From MongoDB to GeminiDB Mongo](#)

DDS Serving as the Source

- [From DDS to MongoDB](#)

Redis Serving as the Source

- [From Redis to GeminiDB Redis](#)

Redis Cluster Serving as the Source

- [From Redis Cluster to GeminiDB Redis](#)

GeminiDB Redis Serving as the Source

- [From GeminiDB Redis to Redis](#)
- [From GeminiDB Redis to Redis Cluster](#)

13.2 Constraints on Synchronization Tasks

To ensure that data synchronization tasks run properly, DRS has certain constraints on different data flow scenarios.

MySQL Serving as the Source

To the Cloud

- [From MySQL to MySQL](#)
- [From MySQL to TaurusDB](#)
- [From MySQL to PostgreSQL](#)
- [From MySQL to GaussDB Centralized](#)
- [From MySQL to GaussDB Distributed](#)

From the Cloud

- [From MySQL to MySQL](#)
- [From MySQL to Kafka](#)
- [From MySQL to Oracle](#)

Self-built -> Self-built

- [From MySQL to Kafka](#)
- [From MySQL to GaussDB Centralized](#)

PostgreSQL Serving as the Source

To the Cloud

- [From PostgreSQL to PostgreSQL](#)
- [From PostgreSQL to GaussDB Centralized](#)
- [From PostgreSQL to GaussDB Distributed](#)

From the Cloud

- [From PostgreSQL to PostgreSQL](#)
- [From PostgreSQL to Kafka](#)

Self-built -> Self-built

- [From PostgreSQL to Kafka](#)

Oracle Serving as the Source

To the Cloud

- [From Oracle to MySQL](#)

- [From Oracle to TaurusDB](#)
- [From Oracle to PostgreSQL](#)
- [From Oracle to GaussDB Centralized](#)
- [From Oracle to GaussDB Distributed](#)
- [From Oracle to DDM](#)

Self-built -> Self-built

- [From Oracle to Kafka](#)
- [From Oracle to GaussDB Centralized](#)
- [From Oracle to GaussDB Distributed](#)

DDM Serving as the Source

To the Cloud

- [From DDM to MySQL](#)
- [From DDM to GaussDB Distributed](#)
- [From DDM to GaussDB Centralized](#)
- [From DDM to DDM](#)

From the Cloud

- [From DDM to MySQL](#)
- [From DDM to Oracle](#)
- [From DDM to Kafka](#)

DB2 for LUW Serving as the Source

To the Cloud

- [From DB2 for LUW to GaussDB Centralized](#)
- [From DB2 for LUW to GaussDB Distributed](#)

Self-built -> Self-built

- [From DB2 for LUW to GaussDB Centralized](#)
- [From DB2 for LUW to GaussDB Distributed](#)

TiDB Serving as the Source

To the Cloud

- [From TiDB to TaurusDB](#)

Microsoft SQL Server Serving as the Source

To the Cloud

- [From Microsoft SQL Server to GaussDB Centralized](#)
- [From Microsoft SQL Server to GaussDB Distributed](#)

- [From Microsoft SQL Server to Microsoft SQL Server](#)

MongoDB Serving as the Source

To the Cloud

- [From MongoDB to DDS](#)

DDS Serving as the Source

From the Cloud

- [From DDS to MongoDB](#)
- [From DDS to Kafka](#)

MariaDB Serving as the Source

To the Cloud

- [From MariaDB to MariaDB](#)

From the Cloud

- [From MariaDB to MariaDB](#)

TaurusDB Serving as the Source

To the Cloud

- [From TaurusDB to TaurusDB](#)

From the Cloud

- [From TaurusDB to MySQL](#)
- [From TaurusDB to Kafka](#)
- [From TaurusDB to Oracle](#)

GaussDB Centralized Serving as the Source

From the Cloud

- [From GaussDB Centralized to MySQL](#)
- [From GaussDB Centralized to Oracle](#)
- [From GaussDB Centralized to Kafka](#)
- [From GaussDB Centralized to GaussDB Distributed](#)
- [From GaussDB Centralized to GaussDB Centralized](#)

Self-built -> Self-built

- [From GaussDB Centralized to Oracle](#)
- [From GaussDB Centralized to Kafka](#)
- [From GaussDB Centralized to GaussDB Centralized](#)

GaussDB Distributed Serving as the Source

From the Cloud

- [From GaussDB Distributed to MySQL](#)
- [From GaussDB Distributed to Oracle](#)
- [From GaussDB Distributed to Kafka](#)
- [From GaussDB Distributed to GaussDB Distributed](#)
- [From GaussDB Distributed to GaussDB Centralized](#)

Self-built -> Self-built

- [From GaussDB Distributed to Oracle](#)
- [From GaussDB Distributed to Kafka](#)
- [From GaussDB Distributed to GaussDB Distributed](#)

13.3 Constraints on DR Tasks

To ensure that data disaster recovery tasks run properly, DRS has certain constraints on different data flow scenarios.

- [From MySQL to MySQL \(Single-Active DR\)](#)
- [From MySQL to TaurusDB \(Single-Active DR\)](#)
- [From DDM to DDM \(Single-Active DR\)](#)
- [From TaurusDB to TaurusDB \(Single-Active DR\)](#)
- [From MySQL to MySQL \(Dual-Active DR\)](#)
- [From TaurusDB to TaurusDB \(Dual-Active DR\)](#)

14 Accessing DRS

Procedure

If you have not registered a Huawei ID and enabled Huawei Cloud services, follow the instructions provided in [Registering a HUAWEI ID and Enabling HUAWEI CLOUD Services](#) to register an account at Huawei Cloud official website. After the registration is successful, you can access all Huawei Cloud services, including DRS, RDS, and DDS.

If you have registered a Huawei ID and enabled Huawei Cloud services, you can log in to the management console and [access your DRS](#).

15 Related Services

RDS

DRS can migrate data from your databases to the RDS databases in the cloud. For more information, see [Relational Database Service User Guide](#).

Supported network types during migration to RDS:

- VPC
- VPN
- Direct Connect
- Public network

DDS

DRS can migrate data from your databases to the DDS databases in the cloud. For more information about DDS, see [Document Database Service User Guide](#).

Supported network types during migration from MongoDB databases to DDS:

- VPC
- VPN
- Direct Connect
- Public network

DDM

DRS helps you migrate data from your databases to Distributed Database Middleware (DDM) on the cloud. For more information about DDM, see [Distributed Database Middleware User Guide](#).

Supported network types during migration to DDM:

- VPC
- VPN
- Direct Connect
- Public network

TaurusDB

DRS can migrate data from your databases to TaurusDB on the current cloud. For more information about TaurusDB, see [TaurusDB User Guide](#).

Supported network types during migration to TaurusDB on the current cloud:

- VPC
- VPN
- Direct Connect
- Public network

GaussDB

DRS can migrate data from your databases to GaussDB on the current cloud. For more information about GaussDB, see [GaussDB User Guide](#).

Supported network types during synchronization to GaussDB on the current cloud:

- VPC
- VPN
- Direct Connect
- Public network

IAM

Identity and Access Management (IAM) manages permissions for DRS.

Only users with the DRS administrator permissions can use DRS. To apply for DRS administrator permissions, you can contact the security administrator or apply for a user with DRS administrator permissions.

For more information about IAM, see [Identity and Access Management User Guide](#).

CTS

Cloud Trace Service (CTS) provides records of operations on cloud service resources, enabling you to query, audit, and backtrack operations.

For more information about CTS, see [Cloud Trace Service User Guide](#).

Cloud Eye

Cloud Eye is an open monitoring platform that helps you monitor DRS resources in real time. It reports alarms promptly to ensure that services are running properly.

For more information about Cloud Eye, see [Cloud Eye User Guide](#).

OBS

Object Storage Service (OBS) provides data storage capabilities that are massive, secure, highly reliable, and low-cost.

For more information about OBS, see [OBS Browser+ Operation Guide](#).

DCC

Dedicated Computing Cluster (DCC) provides dedicated, physically isolated computing resource pools on Huawei Cloud, allowing you to use physical computing devices and resources exclusively.

DRS can use physical resources provided by DCC to create instances to migrate or synchronize data in real time.

For more information about DCC, see [Dedicated Computing Cluster User Guide](#).

Simple Message Notification

Simple Message Notification (SMN) can push notifications based on Huawei Cloud users' requirements. End users can receive notifications through HTTP, HTTPS, and applications. You can also integrate application functions through SMN to reduce system complexity.

For more information about SMN, see [Simple Message Notification User Guide](#).

Enterprise Management

You can create enterprise projects based on the enterprise organization structure. Then you can manage resources across different regions by enterprise project, grant different permissions to user groups, and add them to enterprise projects.

For more information about Enterprise Management, see [Enterprise Management User Guide](#).

16 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.

Currently, only IPsec VPN is supported.

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.

Synchronization Instance

A synchronization instance refers to an instance that facilitates the synchronization process. It exists in the whole lifecycle of a synchronization task.

DRS uses the synchronization instance to connect to the source database, read source data, and synchronize the data to the destination database.

Synchronization Log

A synchronization log refers to the log generated during database synchronization. Synchronization 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. The task cannot start until all check items are successful.

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

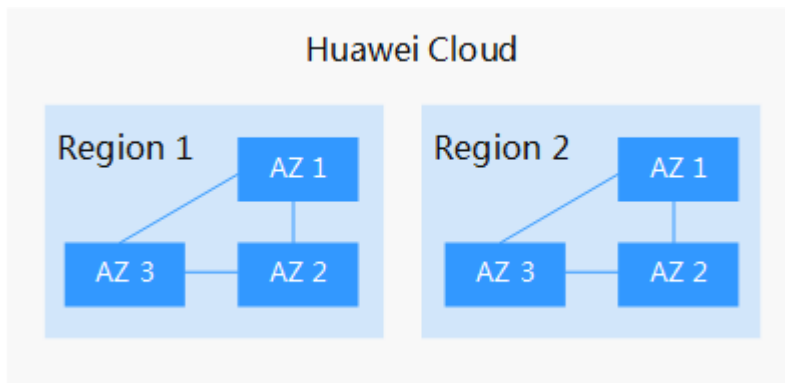
DRS requires that either the source or destination database is on the current cloud. **Out of the cloud** means that the source database must be on the current cloud.

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.

- Regions are divided from the dimensions of geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified as universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides services of the same type only or for specific tenants.
- An AZ contains one or multiple physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to allow you to build cross-AZ high-availability systems.

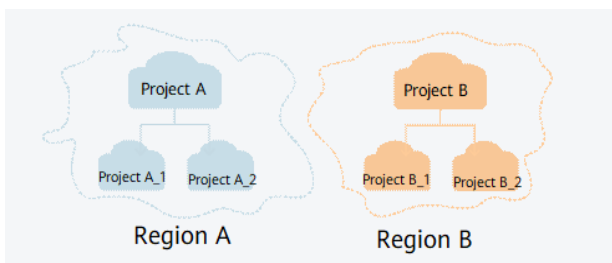
Figure 16-1 shows the relationship between regions and AZs.

Figure 16-1 Region and AZ

Huawei Cloud provides services in many regions around the world. You can select a region and AZ as needed. For more information, see [Huawei Cloud Global Regions](#).

Project

A project corresponds to a region. Projects group and isolate resources (including compute, storage, and network resources) across physical regions. Users can be granted permissions in a default project to access all resources in the region associated with the project. If you need more refined access control, create subprojects under a default project and create resources in subprojects. Then you can assign users the permissions required to access only the resources in the specific subprojects.

Figure 16-2 Project isolating model

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 for MySQL or TaurusDB instances, DRS automatically creates temporary accounts **drsFull** and **drsIncremental** for the destination database. DRS uses these accounts to migrate data. After the task is complete, DRS automatically deletes them.

NOTICE

- Do not use the drsFull and drsIncremental accounts created by users as the account for database connections for a DRS task.
 - 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 or a synchronization instance fails, it automatically fails over to the standby host, preventing service interruption and improving the success rate of migration.

If a replication or synchronization 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 or synchronization 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
- Full synchronization
- Incremental synchronization