

Data Transmission Service

Quick Start

Quick Start

Data migration is easy to use. Only three steps are required to configure a migration task.

Below takes data migration from MySQL to RDS for MySQL as an example to guide through configuring migration task. Other storage engines can follow similar procedures.

Prerequisites

Create an RDS database

During data migration, if the database name does not comply with **RDS-defined Specifications** (a maximum of 64 characters, starting with a letter and ending with a digit. Lowercase letter, digit, underscore and strike-through are applicable). you need to create the database on the RDS console before configuring the migration task.

Create a migration account

To create a migration task, you need to provide the migration accounts of the source instance and the target RDS instance. For the database permissions required by various storage engines, see descriptions in **Product Manual**.

Operation procedure

After you create the RDS databases and migration accounts, you can start to configure a migration task. Follow the below steps to configure a migration task:

1. Log in to DTS Console, click **Create Migration Task** on the top right corner.
2. **Enter the task name and configure the source database and the target database.**

1. Source endpoint and target endpoint 2. Migration class and list 3. Pre-check

* Task name: jiangluteest

Source library

* Instance type: On-premises databases

* Instance region: East China 1

* Database engine: MySQL

* Host name or IP address: rm-bp16w2438yct639x9.mysql.rds.aliyuncs.com

* Port: 3306

* Database account: jiangluteest

* Database password:

Target library

* Instance type: RDS instance

* Instance region: North China 2

* RDS instance ID: rm-dj134k91127zz5293

* Database account: jiangluteest

* Database password:

Cancel Migration evaluation Authorize whitelist and enter into next step

In this step, we will mainly configure the migration task name, the source database and the target database. Where:

1) Task name

By default, DTS automatically generates a name for the migration task. You can edit the name to reflect its special services.

2) Source database information

Instance type: **On-permission databases**

Instance Region: Select the region closest to the source instance.

Database type: **MySQL**

Host name or IP address: **MySQL Instance Connection String**

Port: **the listener port of MySQL instance**

Account: **The access account to the MySQL instance**

Password: **The access account password of the MySQL instance**

3) Target database information

Instance type: **RDS Instance**

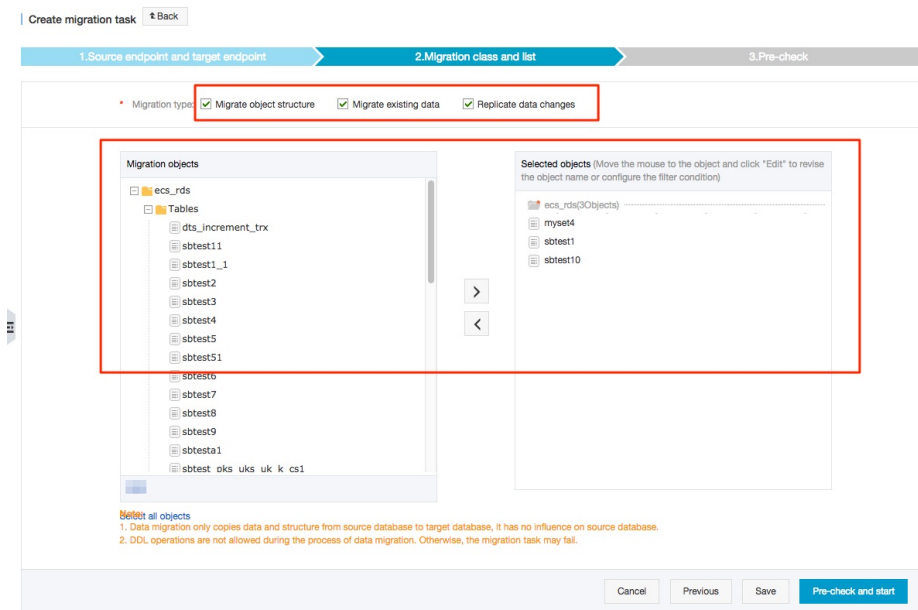
Instance region: The region of the RDS instance.

RDS instance ID: Select the ID of the target RDS instance for migration

Database account: The access account to the RDS instance

Database password: The password of the above database account

3. Configure migration object and migration type



In this step, you need to configure the migration object and migration type. Where:

1) Migration type

The migration type includes migrating object structure, migrating existing data and migrating data changes.

For migrating existing data, you should select migrating schema structure and migrating existing data.

For an zero downtime migration, you should select migrating schema structure, migrating existing data and migrating change data.

2) Migration objects

The migration object can be a database, a table or a column.

By default, the object names should be the same as those in the source database after they are migrated to the target database. Otherwise, you need to use the object name mapping function provided by DTS. Details can be found in [Database Table Column Mapping](#).

4.Pre-check

A pre-check will be performed before a migration task is formally started. If the pre-check fails, click **Failure** to see the details of failed items, solve the problems and perform pre-check again.

Pre-check

Pre-check failed 90%

Check item	Check content	Check result
Check database availability	Check whether the database for target database to be migrated in is available	Success
Check source database permission	Check whether account permissions for the source database meet the requirements for migration	Success
Check target database permission	Check whether account permissions for the target database meet the requirements for migration	Success
Check objects with the same name	Check whether there are any structure objects having the same names with objects to be migrated in the target database	Failed

Cancel

Click **View** to see the reason and solution.

View details

Check content

Check item:Check objects with the same name

Check content:Check whether there are any structure objects having the same names with objects to be migrated in the target database

Check result

Check result:Failed

Failure reason:

There are structure objects in the target database with the same names as the objects to be migrated

Table dts.test4 CHECK__ERROR_SAME_OBJ1 Table

Table dts.test2 CHECK__ERROR_SAME_OBJ1 Table

Table dts.test3 CHECK__ERROR_SAME_OBJ1 Table

Table dts.test CHECK__ERROR_SAME_OBJ1 Table

Table dts.test1 CHECK__ERROR_SAME_OBJ1 Table

Solution:

You may use any of the following three repair modes and try again: (1) Change the name of the object to be migrated in the target database and apply the rename function selected by the migration list (2) Delete or rename objects in the target database (3) Temporarily not migrate objects with the same

Close

After the problem is solved, select the task in the task list, and then re-start pre-check.

5.Start the migration task

After pre-check succeeds, you can start the migration task. Once started, you can view the specific migration status and progress in the task list.

Real time data synchronization is easy to use. Only three steps are required to configure synchronization instance. This section introduces how to use DTS to quickly create a synchronization job between two RDS (MySQL) instances for real time data replication.

Synchronization restrictions

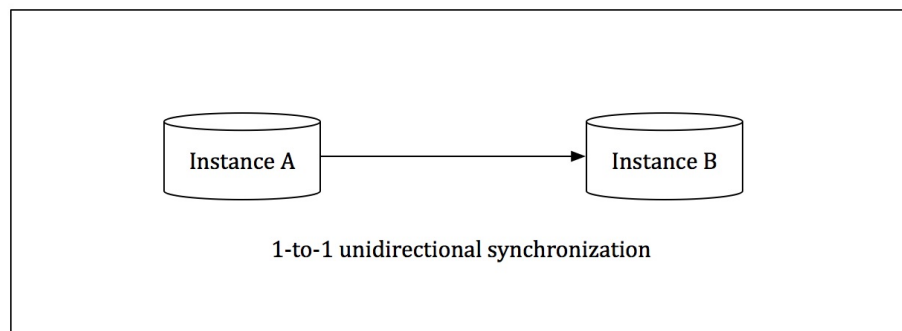
Data source

- At present, real time synchronization only supports RDS MySQL instances.
- The target instance does not support RDS instances in standard access mode and with internet connection addresses only.
- Real time synchronization for RDS instances is not supported in Hong Kong Zone A.


Synchronization architecture

At present, real time synchronization only supports three types of synchronization architecture as below:


- A->B, that is unidirectional synchronization between two instances. And the synchronized object on B must be read-only, otherwise it may cause synchronization instance exceptions.



- A->B/C/D, that is 1-to-N distributed synchronization architecture. The number of target RDS instances is not limited, but the synchronized objects on the target instance must be read-only. Otherwise it may cause synchronization instance exceptions.

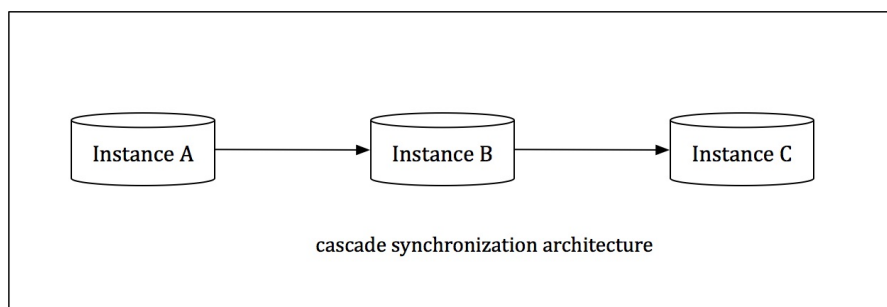


- B/C/D->A, that is N-to-1 data converging architecture. For N-to-1 architecture, the synchronized objects for every synchronization instance must be different to ensure full synchronization.




The following architecture is not supported:

- A->B->C, that is, cascade architecture.



- A->B->A, that is, bidirectional synchronization between Instance A and Instance B.



If an unsupported architecture is set when you configure a synchronization instance, checking **Complex Topology** will fail in pre-check.

Incompatible trigger

When the object to be synchronized is an entire database and the database contains a trigger to update the content of synchronizing table, synchronization data may be inconsistent.

If there is a trigger on A. The body of the trigger is After you insert a row of data into Table A, It will insert a row of data into Table B. In this case, during data synchronization, if you insert a row of data into table A, there will be two row inserted into table B in the target instance .

To solve this problem, you can delete the trigger in the target instance. Data change in Table B is synchronized by DTS.

Operation procedure

Below shows how to create a synchronization job between two RDS instances.

1. Purchase a synchronization instance

- (1) Log in to DTS Console and enter the data synchronization page.
- (2) Click **Create Synchronization Task** on the top right of the console.
- (3) Purchase a synchronization instance before configurations. Synchronization instances now support subscription and pay-as-you-go payment modes. You can select the one you need.

Where:

- Source region

The source region is the region where the source RDS instance is located.

- Target region

The target region is the region where the target RDS instance is located.

- Instance specification

The instance specification affects the performance of synchronization instance. The relationship

between instance type and performance can be found in [Data Synchronization Specifications](#).

- Quantity

The quantity is the number of synchronization instances purchased a time. If you have a pay-as-you-go instance, you can purchase a maximum of 99 instances at a time.

After the synchronization instance is purchased, return to the DTS Console, click **Configure Synchronization instance** on the right of the newly purchased instance to start configurations.

2. Configure synchronization instance connection information

Where:

- Synchronization job name

The synchronization job name is not required to be unique. A name indicating specific services is recommended to facilitate instance identification and management.

- RDS instance ID for synchronization instance

After completing configuration, click **Authorize Whitelist and enter into next step**.

3. Update whitelist for the RDS instance

Add the IP address of the DTS server to the whitelist of the RDS instance for synchronization. This way, you can avoid a failure to create a synchronization job because the DTS server fails to connect to the RDS because of the RDS whitelist.

To ensure the synchronization job stability, do not delete these server IP addresses from the whitelist of RDS instances.

After whitelist authorization, click **Next** to create a synchronization account.

4. Create a synchronization account on the target database

Create a synchronization account on the target RDS instance named **dtssyncwriter**. You cannot delete this account during synchronization. Otherwise, it may cause synchronization job interruption.

5. Select a synchronization object

After a synchronization account is created on the target RDS instance, you can select the synchronization object. The granularity of objects to be synchronized in real time can be subdivided into tables, that is, you can select to synchronize some databases or several tables.

If you select to synchronize the entire database, all the DDL operations in the database (such as create table and drop view) will be synchronized to the target database.

If you select to synchronize a table, only drop/alter/truncate/rename table and create/drop index operations in these tables will be synchronized to the target database.

Note: the rename table operation may cause inconsistent synchronization data. For example, if the synchronization object only includes Table A, instead of Table B and if the source instance executes

rename A to B operation during synchronization, the operations to the renamed Table B won't be synchronized to the target database. To solve this problem, you can choose to synchronize the entire database of Table A and Table B.

After the synchronization object is selected, you can configure the synchronization initialization.

6. Configure synchronize initialization

Initialization is the first step to start a synchronization instance. It will initialize the structure and data of the existing synchronization objects in the source instance on the target instance, use them as the baseline data for subsequent data replication.

Synchronization initialization supports two modes: structural initialization and data initialization. By default, you need to select structural initialization and data initialization.

7. Pre-check

Once all above configuration are completed, pre-check is initiated.

After a synchronization job is configured, DTS checks the limited items. After pre-check succeeds, click **Start** to start the synchronization task.

After a synchronization job started, the synchronization job list is displayed. The just-started task is in **Synchronization Initialization** state. The initialization period depends on the size of the synchronized objects on the source instance. When the initialization is completed, the synchronization instance enters the **Synchronizing** state and a data synchronization is established between the source and target instances.

DTS provides a data subscription function to subscribe to RDS binlog in real time, enabling such scenarios as **lightweight cache update**, **message notify**, **real-time data synchronization with ETL logic**.

To subscribe to RDS binlog in real time, follow these two steps:

1. On the DTS Console, create the subscription instance for RDS instances.
2. Use the SDK provided by DTS to access the subscription instance for real-time subscription and utilize binlog.

Below shows how to create a subscription instance on the DTS console. It is easy to configure, and only three steps are needed. For the management of subscription instances and the use of SDK, see DTS user guid.

Data subscription limits

- Only RDS For MySQL instance is supported.
- The binlog_row_image format of MySQL5.6 binlog must be **Full**.
- MySQL storage engine only supports myisam and innodb.
- MySQL character set only supports latin1, GBK, UTF-8, UTF8mb4 and binary.

Create a subscription instance

TO create a subscription channel, follow the below steps:

1. Purchase a subscription instance

Log in to DTS Console, select **Data Subscription** on the left to enter the data subscription page. Click **Create Data Subscription** on the top right corner to enter the DTS purchase page.

Select a payment mode: Subscription or pay-as-you-go.

Configure below parameters:

- Source instance region

Region of the RDS instance to subscribe to. You cannot change the region after the RDS instance is subscribed.

- Quantity

Number of subscription instances purchased at a time. If you have a pay-as-you-go instance, you can purchase a maximum of 99 instances at a time.

After a subscription instance is purchased, return to the DTS Console and click **Configure subscription channel** on the right of the newly purchased instance to start instance configurations.

2. Configure RDS instance ID

Create data subscription [Back](#)

1. Select instance 2. Select subscription object 3. Pre-check

Subscription name: 1uqf4sktag3(singapore)

Instance type: RDS (MySQL)

Instance region: Ap-southeast

* RDS instance ID: rm-gs59ootxv4j7u05u4

Note: The data subscription in the read-only instance and temporary instance is not supported for data transmission.

Cancel Authorize whitelist and enter into next step

Configure the subscription instance name and ID of the RDS instance to be subscribed, where:

- **Subscription instance name:** An alias of the subscription instance. It has not to be unique. By default, DTS automatically generates a name for each subscription instance. You can edit it to reflect specific services.

- **RDS instance ID:** ID of the RDS instance corresponding to the subscribed binlog.

After configuration, click **Authorize[A2] Whitelist and enter into the next step**.

3. Update the whitelist of the RDS instance.

Add the IP address of the DTS server to the white list of the subscribed RDS instance so as to avoid a failure in creating a subscription instance.

To ensure the stability of a subscription instance, do not delete the IP addresses of DTS server from the whitelist of RDS instances.

4. Select a subscription object

Configure the subscription data type and subscription object, where:

- **Subscription data type:** **Data Change** and **Structure Change**.

Data change means DML (for instance insert/delete/update) and Structure change means DDL (For example create/drop/alter table).

If you subscribe to structure change, DTS will pull all the structure changes on the entire RDS instance. If you only need part of structure changes, you need to filter them at SDK consumption.

- **Subscription object:** **Database** and **Table**. You can subscribe to some databases or several tables.

Once above configuration is completed, you will enter the pre-check stage before starting the subscription instance.

5. Pre-check

After the subscription instance is set up, pre-check starts. After pre-check is passed, you can click **close** on the bottom right corner to start the subscription instance.

6. Start the subscription instance

After the subscription instance is created, it will take about 1 minute to initiate subscription. After initialization, you can **View subscription data**, or **Subscribe to change Data Using SDK**.

So far, the configuration for data subscription instance is completed.