

Data Transmission Service

Quick Start

Quick Start

Purchase DTS instances

This document describes how to purchase different kinds of DTS instances in the DTS console. If you are familiar with the procedure, you can go to the [Data Transmission Service purchase page](#) and place an order directly.

Procedure

Log on to the Alibaba Cloud DTS console.

On the **Overview** page, select the region in which to run the DTS instance, and then click the corresponding shopping cart icon.

On the **Data Transmission Service** purchase page, select the billing method (**Pay-As-You-Go** or **Subscription**).

Note: Data Migration is only available in Pay-As-You-Go. Data Synchronization and Data Subscription are available in both methods.

Complete the **Basic Configuration** and **Purchase Plan**. The price is calculated automatically and displayed on the right side.

For **data migration instances**, the parameter description is as follows:

Parameter	Description
Instance Type	Select Data Migration .
Region	Specify the target instance region with caution. After you create a data migration instance, you cannot change its target region. The existing data migration is free. Only running incremental replication

	incurs configuration fee.
Specification	Select from micro , small , medium , and large . The selected specification affects the data replication performance. For more information about the relationship, see Data migration specifications .
Quantity	Specify how many instances you want to purchase in this order.

For **data subscription instances**, the parameter description is as follows:

Parameter	Description
Instance Type	Select Data Subscription .
Region	Specify the source RDS instance region with caution. After you create a data subscription instance, you cannot change its source region. China (Hong Kong) is only available in Zone B.
Duration	Specify how long you want to subscribe to the data subscription instance. It is only required by the Subscription billing method.
Quantity	Specify how many instances you want to purchase in this order.

For **data synchronization instances**, the parameter description is as follows:

Parameter	Description
Instance Type	Select Data Synchronization .
Source Region	Specify the source instance region with caution. After you create a data synchronization instance, you cannot change its source region. China (Hong Kong) is only available in Zone B.
Target Region	Specify the target instance region with caution. After you create a data synchronization instance, you cannot change its target region. China (Hong Kong) is only available in Zone B.

Specification	Select from micro , small , medium , and large . The selected specification affects the synchronization performance. For more information about the relationship, see Data synchronization specifications .
Duration	Specify how long you want to subscribe to the data synchronization instance. It is only required by the Subscription billing method.
Quantity	Specify how many instances you want to purchase in this order.

Click **Buy Now**.

Confirm your order and complete the payment:

- For a Subscription instance, click **Pay**.
- For a Pay-As-You-Go instance, click **Activate**.

Note: The cross-country synchronization instance is unavailable by default. If you want to purchase a cross-country synchronization instance, submit a ticket or contact your customer manager. When your application passes the back-end review, the purchase entry is included.

Create a data migration task

This document describes how to perform data migration in the DTS console. An example of data migration from MySQL to RDS for MySQL is used to explain the procedure. Other storage engines follow a similar procedure.

Prerequisites

Create a target RDS database.

During data migration, if the database to be migrated is not defined in the target RDS instance, DTS automatically creates it. In some cases where the database name does not comply with the RDS Definition Standard, you must create the target RDS database in the RDS console before configuring the migration task.

The RDS Definition Standard states the RDS database name can be up to 64 characters in length and consist of lowercase letters, numbers, underscore (_), or hyphen (-). It must start with a letter and end with a letter or number.

Create migration accounts.

To create a migration task, you must provide the migration accounts of the source instance and the target RDS instance. For different storage engines, specific database access privileges are also required. For more information, see the corresponding documents in the DTS User Guide. For example,

- Migrate data from on-premises Oracle to RDS for MySQL
- Migrate data from on-premises MySQL to RDS for MySQL

Procedure

Follow these steps to configure a migration task.

Log on to the Alibaba Cloud DTS console.

Click **Data migration** in the left-side navigation pane.

Click **Create migration task**.

Configure the **Source endpoint and target endpoint**.

Task Name: DTS automatically generates a name for the migration task. You can edit the name according to your business requirements.

Source database

Parameter	Description
Instance type	Select On-premises databases .
Instance region	Select the region in closest geographic proximity to the source instance.
Database engine	Select MySQL .
Host name or IP address	Enter the MySQL database connection string.
Port	Select the MySQL database listener port.
Database account	Enter the MySQL database account

	name.
Database password	Enter the MySQL database account password.

Select the **Connection method** (**Non-encrypted connection** or **SSL secure connection**), and then click **Test the connection** to verify the input contents.

Target database

Parameter	Description
Instance type	Select RDS Instance .
Instance region	Select the region of the RDS instance.
RDS instance ID	Select the target RDS instance ID.
Database account	Enter the access account name of the RDS instance database.
Database password	Enter the access account password of the RDS instance database.

Click **Test the connection** to verify the input contents.

Configure **Migration type and migration objects**.

Migration type: includes **Schema migration**, **Existing data migration**, and **Incremental replication**.

- For existing data migration, select schema migration and existing data migration.
- For zero downtime data migration, select schema migration, existing data migration, and incremental replication.

Migration objects: include **database**, **table**, and **column**.

Note: By default, the object names are the same as in the source and target database. To use the object name mapping function, see **Object name mapping**.

Perform **Pre-check** before the migration task is started. If the pre-check fails,

- Click **Failed** to check the details of the failed items.
- Click **View** to view the reason and solution.

Troubleshoot the issues and perform pre-check again until the pre-check is successful.

After the pre-check passes, start the migration task.

When the task is started, you can view the migration status and progress in the task list.

Create a data synchronization task

This document introduces how to create a data synchronization task in the DTS console. An example of building a synchronization link between two RDS for MySQL instances is used to describe the procedure.

Prerequisites

The data synchronization function applies restrictions on instance types, synchronization architectures, and trigger objects. Before starting the configuration, make sure that the following restrictions are well met.

Instance types

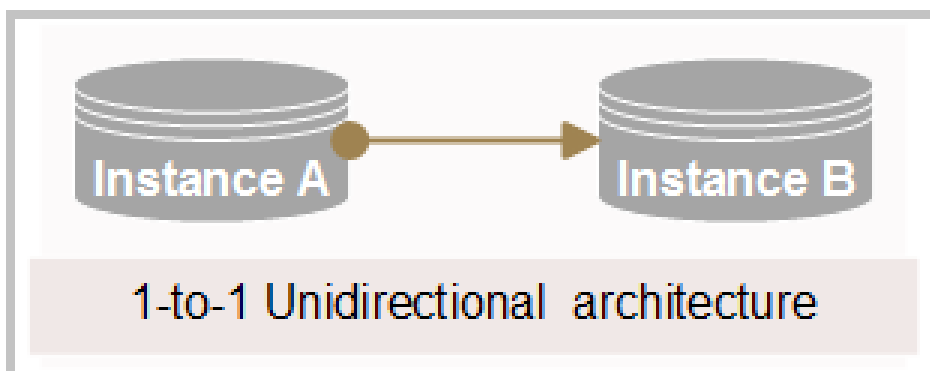
Currently, real-time synchronization is available for **RDS for MySQL** instances.

- The target instance does not support RDS instances that use standard access mode and the Internet addresses.
- RDS instances in China (Hong Kong) Zone A are unavailable.

Synchronization architectures


Real-time synchronization supports the following types of synchronization architectures:

A->B, or unidirectional synchronization between two instances.



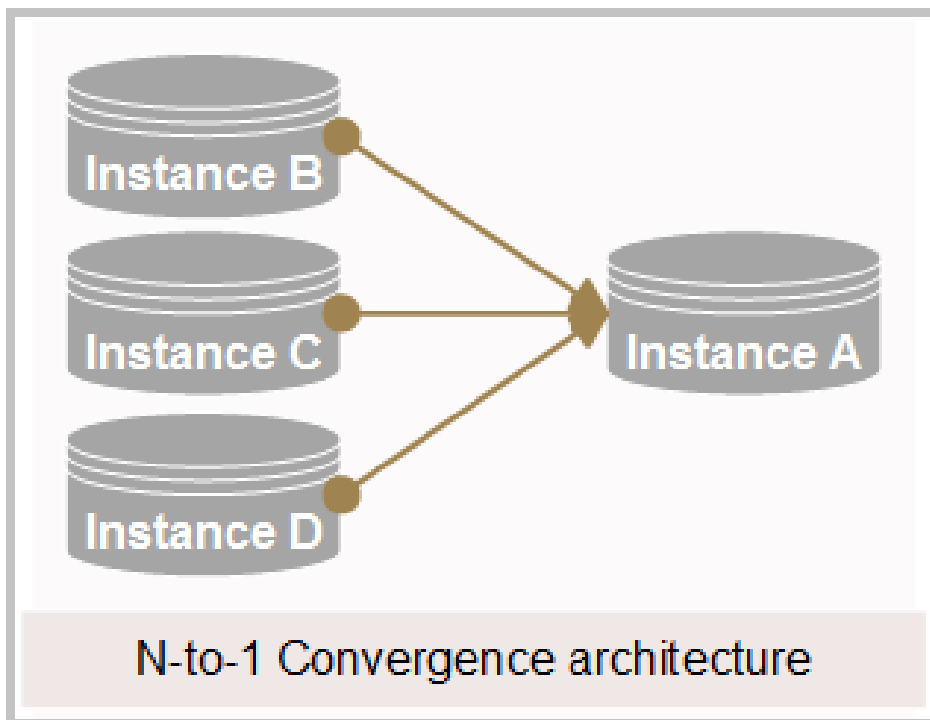
The synchronization objects on B must be granted read-only permission only. Otherwise, it may cause synchronization link exceptions.

A->B/C/D, or 1-to-N distributed synchronization architecture.



The number of target RDS instances can be set on your demand, but the synchronization objects on the target instances must be granted read-only permission only. Otherwise, it may cause synchronization link exceptions.


B/C/D->A, or N-to-1 data convergence architecture.




The synchronization objects on each synchronization instance must be different from each other, ensuring complementary integration.

The following architectures are **NOT** available:

A->B->C, or cascade architecture.



A->B->A, or bidirectional synchronization between Instance A and Instance B.



If unsupported architectures are detected in the synchronization configuration, the pre-check returns a **Complex Topology** error.

Incompatible triggers

For full database synchronization, if the source database contains triggers that can update the content of the synchronized table within the database, data inconsistency may occur between the source and target instances.

Suppose that the source instance contains Table A and Table B. When a row is inserted into Table A, a trigger in Table A inserts a row into Table B. The synchronization task synchronizes the source instance to the target instance, where Table A and Table B are represented as Table A' and Table B'.

During synchronization, the target instance continuously replicates full amount data from the source instance, including Table A and Table B. If you insert a row (1) into Table A, the trigger in A inserts a row (2) into Table B. The row you inserted (1) and the row inserted by the trigger (2) are synchronized respectively to the target instance Table A' and Table B'. When the row you inserted (1) is updated to Table A', the trigger in A' inserts a row (2') in Table B'.

As a result, one row insertion (1) in the source Table A triggers two row insertions (2 and 2') in the target Table B. Therefore, the data between source and target is inconsistent.

To avoid this problem, delete the trigger in the target instance and synchronize Table B from the source instance.

Procedure

Before going ahead, make sure you have at least one synchronization instance available, or you can purchase one on the [Data Transmission Service purchase page](#). For specific instructions, see [Purchase DTS instances](#).

Follow these steps to create a synchronization task in the DTS console.

Log on to the Alibaba Cloud DTS console and enter the **Data Synchronization** page.

Locate the synchronization instance to be configured and click **Configure synchronization instance** in the **Operation** list.

Select the source and target instances. The parameters include:

- **Synchronization task name:** Enter a name for the synchronization task. The name is not required to be unique.
- **Source instance information:** Select the source **Instance ID**, and specify the **Connection method**.
- **Target instance information:** Select the target **Instance ID**, and specify the **Connection method**.

Note:

- The source and target instances must be different instances.
- The connection methods include **Non-encrypted connection** and **SSL secure connection**. The secure connection consumes more CPU resources.

After the configuration, click **Authorize whitelist and enter into next step**.

Wait for the backend process to add the IP address of the DTS server to the specified RDS instance whitelists and create a synchronization account on the target RDS instance. After the process is complete, click **Next**.

- By updating the RDS instance whitelists, the DTS server is granted access permissions. This can avoid connecting failures due to an RDS whitelist restriction. We recommend that you do not remove the server IP addresses from the whitelists of RDS instances during synchronization.
- DTS automatically creates a synchronization account on the target RDS instance,

and names it **dtssyncwriter**. We recommend that you do not remove this account during synchronization. Otherwise, it may cause synchronization task interruption.

Select the synchronization object (database or table). You can specify multiple databases or tables.

- If you specify the entire database, all the DDL operations in the database (such as create table and drop view) are replicated to the target database.
- If you specify a table, only drop/alter/truncate/rename table and create/drop index operations in the table are replicated to the target database.

Note: The renaming table operation may cause inconsistent data after synchronization.

For example, assume that the synchronization objects include Table A, but do not include Table B. If the source instance renames A to B during synchronization, all data changes in the source Table B are not replicated to the target instance. To avoid such a problem, choose to synchronize the entire database which includes both Table A and Table B.

Click **Next**.

Perform **Advanced setup**. DTS creates a replica of synchronization objects on the target instance as the base-line data.

The setup includes two types: schema initialization and existing data initialization. Typically, both of the initialization types are required.

Perform **Pre-check**. DTS examines a predefined checklist for possible causes of failure.

Click **Start** to start the synchronization link. After being started, the synchronization task is added to the synchronization list.

Note: The link stays in **Synchronization Initialization** status for a period of time, depending on the size of the synchronized objects.

When the initialization is complete, the synchronization instance status becomes **Synchronizing**, and the synchronization link is established between the source and target instances.

Configure a data subscription channel

DTS can subscribe to RDS binlog, allowing you to view incremental data in real time and consume the data with built-in SDK. You can use the data subscription function for lightweight cache update, message notification, and real-time data synchronization with ETL logic.

At a high level, when using data subscription, you follow these steps:

- Create a data subscription instance on the Data Transmission Service purchase page.
- Configure a subscription channel to subscribe to an RDS instance in the DTS console.
- Use DTS SDK to access the subscription channel for obtaining and consuming binlog.

This document describes how to configure a subscription channel to subscribe to an RDS instance in the DTS console.

Prerequisites

Data subscription applies the following restrictions:

- The source instance only supports RDS for MySQL.
- For MySQL 5.6, the `binlog_row_image` format of binlog must be **Full**.
- The MySQL storage engine only supports `myisam` and `innodb`.
- The MySQL character set only supports `latin1`, `GBK`, `UTF-8`, `UTF8mb4`, and `binary`.

Procedure

Before going ahead, make sure you have at least an available subscription instance, or you can purchase one through the [Data Transmission Service purchase page](#). For more information, see [Purchase DTS instances](#).

Follow these steps to configure a subscription channel in the DTS console.

Log on to the [Alibaba Cloud DTS console](#).

Click **Data subscription** in the left-side navigation pane.

Click **Configure subscription channel** under the target subscription instance.

In the **Select instance** stage, complete the following configurations, and then click **Authorize whitelist and enter into next step**.

- **Subscription name:** specifies the subscription instance according to your business requirements, or leave it as it is. The name is an alias of the subscription channel and is not required to be unique.

- **Instance type:** specifies **RDS instance**.
- **RDS Instance ID:** specifies the ID of the RDS instance which is subscribed to.

In the **Select subscription object** stage, complete the following configurations, and click **Save and pre-check**.

Data to be subscribed: specifies **Data update** or **Structure update** or both.

- Data update indicates DML (Data Manipulation Language, used to manipulate data itself), such as insert, delete, and update.
- Structure update indicates DDL (Data Definition Language, used to define data structures), such as create table, drop table, and alter table.

If you subscribe to structure change, DTS pulls all the structure changes on the RDS instance. The obtained structure changes can be filtered for SDK consumption as needed.

Subscription objects: specifies multiple databases or tables.

Wait until the pre-check passes, and then click **Close** to start the subscription instance.

After the subscription instance is started, the subscription channel begins initializing. After initialization, you can **View subscription data**, or **Subscribe to change data by using SDK**.