


# **Data Transfer Service**

## **Practical Tutorial**

### **Product Documentation**



## Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

## Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

## Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

# Contents

## Practical Tutorial

- Synchronizing Local Database to the Cloud
- Creating Two-Way Sync Data Structure
- Creating Many-to-One Sync Data Structure
- Creating Multi-Site Active-Active IDC Architecture
- Selecting Data Sync Conflict Resolution Policy
- Using CLB as Proxy for Cross-Account Database Migration
- Migrating Self-built Databases to Tencent Cloud Databases via CCN

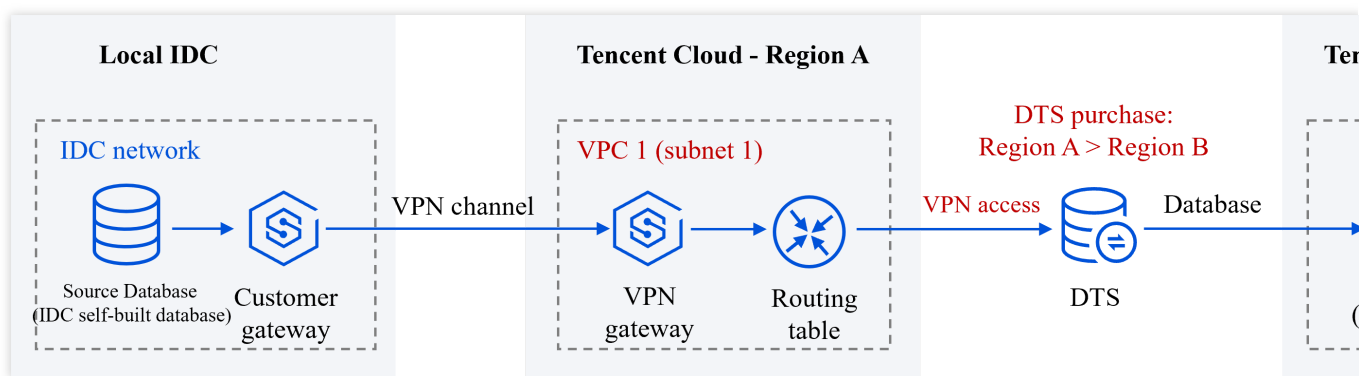
# Practical Tutorial

## Synchronizing Local Database to the Cloud

Last updated : 2024-07-11 15:35:32

### Overviews

This article provides instructions for using DTS to synchronize local IDC self-built database services to Tencent Cloud. In the scenario of service migration, after the local database is synchronized to the cloud, a cutover must be performed. To prevent abnormalities in the cloud database after the service cutover, it is recommended to use the DTS data synchronization module to configure a reverse escape link, so that when the database on the cloud is synchronized abnormally, the service can switch back to the local database.



### Preparations

#### 1. Preparation for Network Connection Establishment

Using DTS for database synchronization requires opening up connectivity between the source/target database and Tencent Cloud VPC, so that DTS can connect to the source/target database.

In this example, the source database is a self-build IDC database. DTS can connect through Public Network/VPN Access/Direct Connect/CCN Methods. We will use VPN Access as an example. The target database is a Tencent Cloud database instance.

1. Connect the local IDC **nearby access** to Tencent Cloud VPC.

For specific network connection operations of VPN Access, see [Interworking Between local IDC and Tencent Cloud](#). If you wish to use other access methods, you can also see [Interworking Between local IDC and Tencent Cloud](#) to view network connection configuration principles.

2. When a DTS task is purchased later, **Source Instance Region** needs to select the region where the source library's Tencent Cloud VPC is located, which is VPC1's region Beijing. **Target Instance Region** needs to select the region where the target database is located, which is Guangzhou.
3. In the subsequent DTS task configuration, in the source library settings, **Access Type** choose VPN Access, for **VPC** and **Subnet** , select VPC1 and choose one of its subnets, subnet1; in the target library settings, **Access Type** choose Database.

## 2. Preparation for Account and Permission


Create an account for executing DTS tasks and grant it permissions. The following is an introduction using MySQL as an example.


Authorization for the source database is as follows:



```
#Creating Execution Task Account
CREATE USER 'account'@'%' IDENTIFIED BY 'password';
#Grant Permissions
GRANT RELOAD, LOCK TABLES, REPLICATION CLIENT, REPLICATION SLAVE, SHOW VIEW, PROCES
GRANT ALL PRIVILEGES ON `__tencentdb__`.* TO 'account'@'%;
FLUSH PRIVILEGES;
```

Authorization for the target database is as follows:





```
#Creating Execution Task Account
CREATE USER 'account'@'%' IDENTIFIED BY 'password';
#Grant Permissions
GRANT ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE
FLUSH PRIVILEGES;
```

## Notes

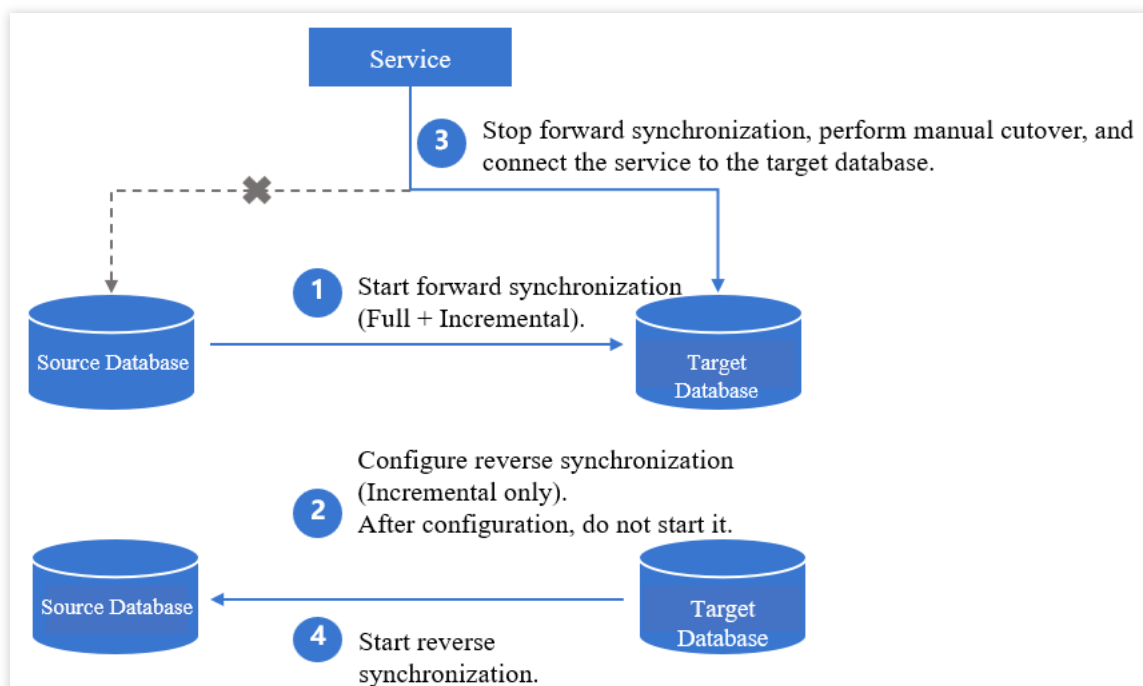


Forward synchronization and reverse synchronization are two independent unidirectional sync tasks. Each independent task's constraints, operational limits, etc., need to meet the basic requirements of the sync tasks, see [Data Synchronization](#) section for corresponding sync links.

DBbridge, when executing full data synchronization, will occupy certain source database resources, which may lead to an increase in the load on the source database, adding to the database's own pressure. If your database configuration is too low, it's recommended to proceed during the business off-peak period.

## Summary of Directions

In scenarios where DTS is used for database migration, to prevent data anomalies in the target database after cutover, it is recommended to use data synchronization configuration to establish a reverse escape route. This allows for a business switchback to the source database if an anomaly occurs in the target database.



1. Configure and initiate the forward synchronization task, opting for full + incremental synchronization.

Key configurations in the forward task: For **Initialization Type**, select Structure Initialization + Full Data Initialization; for **Existing Table with Same Name**, choose Pre-validate and report error.

2. Configure the reverse task for incremental synchronization only. After configuration, do not start it immediately.

Key configurations in the reverse task: For **Initialization Type**, do not select any; for **Existing Table with Same Name**, choose Ignore and Continue Execution.

3. Forward synchronization is completed, stop the forward task, perform manual cutover, and connect the service to the target database.

4. Start the reverse synchronization task to synchronize the incremental data from the target database back to the source database.

5. (Optional) If the data in the target database is abnormal after cutover, stop the reverse synchronization and switch the service back to the source database.

## Detailed Operation Information

### Note:

The overviews for different database links are similar. The following introduces syncing from MySQL to MySQL as an example. For more information, see [Data Synchronization](#) section under Link Configuration Guide.

### Step One: Purchasing DTS

Log in to the data synchronization purchase page, select the appropriate configuration, and click **Buy Now**.

1. Forward task.

**Source Instance Region** is to choose the region associated with the source library's Tencent Cloud VPC, which is Beijing for VPC1. **Target Instance Region** is to choose the region where the target database is located, which is Guangzhou.

2. Reverse task.

The database type and region selection of the source and target instances are opposite.

### Step Two: Creating and Starting Forward Synchronization

1. After successful purchase, return to the [data sync list](#), and you can see the newly created data sync task, click **Configure** in the **Operation** column to enter the sync task configuration page.

2. On the sync task configuration page, configure the source and target instances and their accounts and passwords, test the connectivity, and click **Next**.

1

Set source and target databases

>

2

Set sync options and objects

>

3

Verify task

Task Configuration

Task Name \*

sync-

Running Mode \*

Immediate execution

Scheduled execution

Source Instance Settings

Source Instance Type \*

MySQL

Source Instance Region

South China(Guangzhou)

Service Provider \*

Others

AWS

Alibaba Cloud

Access Type \*

Public Network

Self-Build on CVM

Direct Connect

VPN Access

Database

CCN

Cross-/Intra-Account \*

Intra-account

Cross-account

Instance ID \*

cdb-

Account \*

root

Password \*

.....

Test Connectivity

Test passed

Target Instance Settings

Target Instance Type \*

MySQL

Target Instance Region

South China(Guangzhou)

Access Type \*

Public Network

Self-Build on CVM

Direct Connect

VPN Access

Database

CCN

Instance ID \*

cdb-

Account \*

root

Password \*

.....

Test Connectivity

Test passed

Category	Parameter	Description
Task Configuration	Task Name	DTS will automatically generate a task name, which is customizable.
	Running Mode	Immediate execution and scheduled execution are supported.

Source Instance Settings	Source Instance Type	Select the source instance type selected during purchase, which cannot be changed once configured.
	Source Instance Region	Select the source instance region selected during purchase, which cannot be changed once configured.
	Service Provider	For a self-built database (such as a CVM-based one) or TencentDB database, select <b>Others</b> . For a third-party cloud database, select the corresponding service provider. In this scenario, select <b>Others</b> .
	Access Type	Select a type based on your scenario. In this scenario, select <b>Direct Connect</b> or <b>VPN Access</b> , and you need to configure VPN-IDC interconnection as instructed in <a href="#">Direct Connect or VPN Access: Configuring VPN-IDC Interconnection</a> . For the preparations for different access types, see <a href="#">Overview</a> . Public Network: The source database can be accessed through a public IP. Self-Build on CVM: The source database is deployed in a <a href="#">CVM</a> instance. Direct Connect: The source database can be interconnected with VPCs through <a href="#">Direct Connect</a> . VPN Access: The source database can be interconnected with VPCs through <a href="#">VPN Connections</a> . Database: The source database is a TencentDB instance. CCN: The source database can be interconnected with VPCs through <a href="#">CCN</a> . VPC: The source and target databases are both deployed in Tencent Cloud <a href="#">VPCs</a> . To use the VPC access type, <a href="#">submit a ticket</a> for application.
	VPC-based Direct Connect Gateway/VPN Gateway	Only VPC-based Direct Connect gateway is supported. You need to confirm the network type associated with the gateway. VPN Gateway: Select a VPN Gateway instance.
	VPC	Select a VPC and subnet associated with the VPC-based Direct Connect Gateway or VPN Gateway.
	Host Address	IP address or domain name for accessing the source MySQL instance.
	Port	Port for accessing the source MySQL instance.
	Account	Account of the source instance, which must have the required permissions.
	Password	Password of the source instance account.

Target Instance Settings	Target Instance Type	The target instance type selected during purchase, which cannot be changed.
	Target Instance Region	The target instance region selected during purchase, which cannot be changed.
	Access Type	Select a type based on your scenario. In this scenario, select <b>Database</b> .
	Instance ID	Target instance ID.
	Account	Account of the target instance, which must have the required permissions.
	Password	Password of the target instance account.

3. On the **Set sync options and objects** page, set the data initialization, data sync, and sync object options and click **Save and Go Next**.

**Note:**

If you only select **Full data initialization** for **Initialization Type**, the system will assume by default that you have created the table structures in the target database and will neither sync table structures nor check whether the source and target databases have tables with the same name. Therefore, if you select **Precheck and report error** for **If Target Already Exists**, the precheck and error reporting feature won't take effect.

If you select **Full data initialization** only, you need to create the table structures in the target database in advance. If you want to rename a table (for example, rename table A table B) during sync, you must select the entire database (or entire instance) where table A resides rather than only table A as the **sync object**; otherwise, the system will report an error.

1 Set source and target databases

2 Set sync options and objects

3 Verify task

### Data Initialization Option

Initialization Type

Structure Initialization Full data initialization

If Target Already Exists \*

Precheck and report error Ignore and execute

### Data Sync Option

Primary Key Conflict Resolution \*

Report Ignore Overwrite

SQL Type

DML

Insert Update Delete

DDL

DDL

Custom DDL \*

Database

Create Drop Alter

Table

Create Drop Alter Truncate Rename

View

Create Drop

Index

Create Drop

### Sync Object Option

Advanced Migration Object ⓘ

☒ Procedure
 ☒ Function

Advanced objects can only be copied once, which means you cannot copy new objects once the task is started.

Up to 200 results can be displayed. If the objects you need are not shown in the result list, you can search them by object name.

#### Source Database Object

Search database name, supporting fuzzy match

1 database in total, with 1 displayed

More

db-dst

Refresh Select all Clear

#### Selected Object ⓘ

Globally search for original object names, with fuzzy match support

db-dst (Entire database selected)

Unfold all Fold all Select all Clear Revert to Original Name

Category	Parameter	Description
Data Initialization Option	Initialization Type	<p>Structure initialization: Table structures in the source instance will be initialized into the target instance before the sync task runs.</p> <p>Full data initialization: Data in the source instance will be initialized into the target database before the sync task runs. Both options are selected by default, and you can deselect them as needed.</p>

©2013-2022 Tencent Cloud. All rights reserved.

Page 14 of 60

	If Target Already Exists	Precheck and report error: If a table with the same name exists in both the source and target databases, an error will be reported, and the task will stop. Ignore and execute: Full and incremental data will be directly added to tables in the target instance.
Data Sync Option	Conflict Resolution Method	Report: If a primary key conflict is found during data sync, an error will be reported, and the data sync task will be paused. Ignore: If a primary key conflict is found during data sync, the primary key record in the target database will be retained. Overwrite: If a primary key conflict is found during data sync, the primary key record in the source database will overwrite that in the target database.
	SQL Type	Supported operations include INSERT, UPDATE, DELETE, and DDL. If you select Custom DDL, you can select different DDL statement sync policies as needed. For more information, see <a href="#">Setting SQL Filter Policy</a> .
Sync Object Option	Database and Table Objects of Source Instance	Select the objects to be synced. You can select basic databases, tables, views, procedures, and functions. The sync of advanced objects is a one-time operation: only advanced objects already in the source database before the task start can be synced, while those added to the source database after the task start will not be synced to the target database. For more information, see <a href="#">Syncing Advanced Object</a> .
	Selected Object	Database/Table mapping (renaming) is supported. Hover over a database or table name, click the displayed Edit icon, and enter a new name in the pop-up window. When advanced objects are selected for sync, we recommend you not rename databases/tables; otherwise, sync of the advanced objects may fail. Online DDL temp tables can be synced (through tools such as gh-ost or pt-online-schema-change). Click Edit of a table and select a temp table name in the pop-up window. For more information, see <a href="#">Syncing Online DDL Temp Table</a> .

4. On the Verify task page, complete the verification. After all check items are passed, click **Start Task**.

If the verification fails, fix the problem as instructed in [Check Item Overview](#) and initiate the verification again.

**Failed:** It indicates that a check item fails and the task is blocked. You need to fix the problem and run the verification task again.

**Alarm:** It indicates that a check item doesn't completely meet the requirements, and the task can be continued, but the business will be affected. You need to assess whether to ignore the alarm or fix the problem and continue the task based on the alarm message.

Create Verification Task

Query Verification Result

connect db check	Pass
necessary check	Pass
version check	Pass
source instance privilege check	Alarm
simple instance param check	Pass
target instance privilege check	Pass
check if target instance has conflict content	Pass
check if there's enough space in target instance	Pass
source instance binlog param check	Pass
foreign key constraint check	Pass
partial table foreign key constraint check	Pass
view check	Pass
warning param check	Alarm

Previous

Start Task

5. Return to the data sync task list, and you can see that the task has entered the **Running** status.

### Note:

You can click **More** > **Stop** in the Operation column to stop a sync task. You need to ensure that data sync has been completed before stopping the task.

<input type="checkbox"/>	replicate- replicate-	(3 / 3)	Immediate execution	Medium	Monthly subscription Expire at 2022-05-12 11:00:01	MySQL -> MySQL	MySQL	MySQL	Database	Database
		Status: <b>Running</b> Start: 2021-12-16 11:02:07 End: -- Data lag between target and source databases: 0 MB Time lag between target and source databases: 0s								

6. (Optional) You can click a task name to enter the task details page and view the task initialization status and monitoring data.

## Step Three: Configuring Reverse Synchronization

The operations for reverse synchronization are basically the same as those for forward synchronization; only the differences are described below.

1. Set the synchronization source and target database.

In this step, the source and target databases are swapped compared to the forward task.

2. Set synchronization options and synchronization objects.

Initialization Type: Do not select any.

Existing Table with Same Name: Choose Ignore and Continue Execution.

Primary Key Conflict Resolution Mechanism: Choose based on the business scenario.

Synchronization Operation Type: Keep consistent with the forward task.

3. On the validation task page, perform the validation. After passing the validation, start the task.

## Step Four: Service Cutover

Wait until both the data gap and latency gap in the forward synchronization task are 0, then you can start the cutover.



1. After the data validation is correct, proceed with the business cutover. For data validation here, you can use [Creating Data Consistency Check Task \(MySQL\)](#) for auxiliary verification.
2. Terminate the forward synchronization task.
3. Manually connect the source database business to the target database.

### Step Five: Starting Reverse Synchronization

Start the reverse task.

### Step Six (Optional): Reverse Cutover

If data exception of target database is detected, stop reverse synchronization and switch the service back to the source database.

# Creating Two-Way Sync Data Structure

Last updated : 2024-07-08 15:54:46

## Overview

DTS supports two-way data sync between two databases, which can be applied to multi-site active-active scenarios. In a two-way sync task, two one-way sync tasks are created to establish a two-way topology, and data can be written into both database instances at the same time during sync.

Two-way data sync must follow restrictions on one-way sync and relevant operations. For more information, see the appropriate sync scenario in [Databases Supported by Data Sync](#).

## Notes

During full data sync, DTS consumes certain source instance resources, which may increase the load and pressure of the source database. If your database configuration is low, we recommend you sync the data during off-peak hours.

To avoid duplicate data, make sure that the tables to be synced have a primary key or non-null unique key.

You should plan the data in advance. The two source databases are responsible for updating (adding, deleting, and modifying) data with different primary keys so as to avoid problems such as primary key conflict and mutual overwriting of data with the same primary key (for example, data records with primary keys `1`, `3`, and `5` are updated in database A, while data records with primary keys `2`, `4`, and `6` are updated in database B). If there are duplicate primary keys in the two source databases for business reasons, select an appropriate conflict resolution policy as instructed in [Recommended Configurations for Typical Use Cases](#) to make the sync behavior and data meet the expectations.

Prepare the target database and [grant the account executing the sync task the permissions of the source and target databases](#).

## Use Limits

DDL statements can be executed in at most one direction during two-way sync, as the sync linkage should not form a ring (you can run DDL statements in either the forward or reverse direction).

All sync links between MySQL, TDSQL-C for MySQL, MariaDB, Percona, and TDSQL for MySQL support two-way sync except when a TDSQL for MySQL instance with the MariaDB kernel is used as the source or target database.

## Recommended Configurations for Typical Use Cases

A two-way sync task consists of two one-way sync tasks to establish a two-way topology. The creation steps for each one-way sync task are similar to those for a general one-way sync task. They differ only in the following sync option settings:

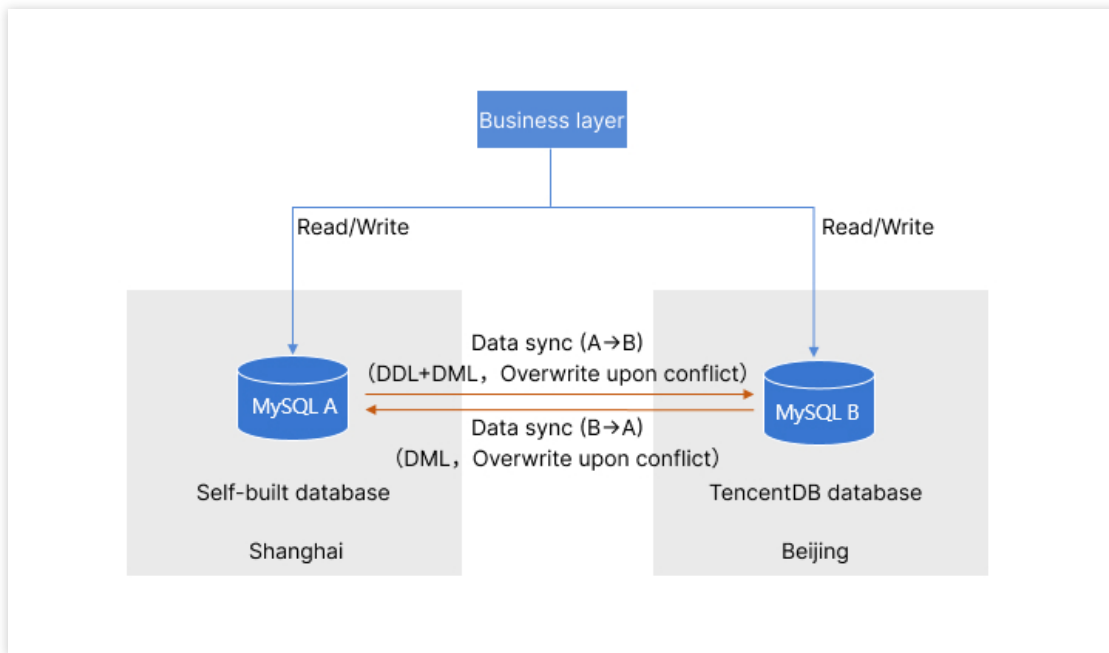
### Sync Option Settings Difference

The following configurations are recommended for typical use cases for your reference.

Scenario	Time Requirements	Sync Task	Initialization Type	If Target Already Exists	Conflict Resolution Method	SQL Type
Scenario 1: Instance A has database/table structures and data, and instance B is empty	Task 2 can be created only after task 1 enters the "incremental sync" phase	Task 1: Forward sync (A < B)	Structure initialization/full data initialization	Precheck and report error	Select an option as needed. Example: If a primary key conflict occurs, and you want the content of database A to prevail, you need to select <b>**Overwrite**</b> for task 1 and <b>**Ignore**</b> or <b>**Report**</b> for task 2. The conflict resolution method takes effect only for the data with primary key conflict.	Select DDL in at most one task. For operation types other than DDL, keep them consistent between the two tasks.
		Task 2: Reverse sync (B > A)	Do not select	Ignore and execute		
Scenario 2: Instance A has database/table structures and data, and instance B has only database/table structures but no data	None	Task 1: Forward sync (A > B)	Full data initialization	Ignore and execute		
		Task 2: Reverse sync (B > A)	Do not select	Ignore and execute		
Scenario 3: Both instances A and B have database/table structures and data	None	Task 1: Forward sync (A > B)	Full data initialization	Ignore and execute		
		Task 2: Reverse sync (B > A)	Full data initialization	Ignore and execute		

## Directions

This document takes creating two-way sync between self-built MySQL database A in Shanghai region and TencentDB for MySQL database B in Beijing region as an example. Initially, A has database/table structures and data, while B is empty. When a primary key conflict occurs, data updates in A shall prevail. For A > B sync, the primary key conflict resolution policy is **Overwrite**, and DDL and DML statements are synced. For B > A sync, the policy is **Ignore**, and only DML statements are synced.



### Creating a sync task 1: Reverse sync (A > B)

1. Log in to the [data sync purchase page](#), select appropriate configuration items, and click **Buy Now**.
2. After successful purchase, return to the [data sync list](#), and you can see the newly created data sync task. You need to configure it before you can use it.
3. In the data sync list, click **Configure** in the **Operation** column to enter the sync task configuration page.
4. On the sync task configuration page, configure the source and target instances and their accounts and passwords, test the connectivity, and click **Next**.

1 Set source and target databases
2 Set sync options and objects
3 Verify task

### Task Configuration

Task Name \*

Running Mode \* Immediate execution Scheduled execution

### Source Instance Settings

Source Instance Type \* MySQL

Source Instance Region South China(Guangzhou)

Service Provider \* Others AWS Alibaba Cloud

Access Type \* Public Network Self-Build on CVM Direct Connect VPN Access Database CCN

Cross-/Intra-Account \* Intra-account Cross-account

Instance ID \*

Account \*

Password \*

Test Connectivity

### Target Instance Settings

Target Instance Type \* MySQL

Target Instance Region South China(Guangzhou)

Access Type \* Public Network Self-Build on CVM Direct Connect VPN Access Database CCN

Instance ID \*

Account \*

Password \*

Test Connectivity

Database A

Database B

Category	Parameter	Description
Task Configuration	Task Name	DTS will automatically generate a task name, which is customizable.
	Running Mode	Immediate execution and scheduled execution are supported.
Source Instance Settings	Source Instance Type	The database A type selected during purchase, which cannot be changed.
	Source Instance Region	The database A region selected during purchase, which cannot be changed.

	Service Provider	Select <b>**Others**</b> .
	Access Type	For a third-party cloud database, you can select <b>**Public Network**</b> generally or select <b>**VPN Access**</b> , <b>**Direct Connect**</b> , or <b>**CCN**</b> based on your actual network conditions. In this scenario, <b>**Public Network**</b> is selected as an example. For the preparations for different access types, see <a href="#">Overview</a> .
Target Instance Settings	Target Instance Type	The target database B type selected during purchase, which cannot be changed.
	Target Instance Region	The target database B region selected during purchase, which cannot be changed.
	Access Type	In this scenario, select <b>**Database**</b> .
	Instance ID	Instance ID of database B.
	Account	Account of database B, which must have the required permissions.
	Password	Password of database B.

5. On the **Set sync options and objects** page, set the data initialization, data sync, and sync object options and click **Save and Go Next**.

1 Set source and target databases > 2 Set sync options and objects > 3 Verify task

### Data Initialization Option

Initialization Type Structure Initialization Full data initialization

If Target Already Exists \* Precheck and report error Ignore and execute

### Data Sync Option

Primary Key Conflict Resolution \* Report Ignore Overwrite

### SQL Type

DML Insert Update Delete

DDL DDL

Custom DDL \* ☐

### Sync Object Option

Advanced Migration Object ⓘ ☒ Procedure ☒ Function

Advanced objects can only be copied once, which means you cannot copy new objects once the task is started.

ⓘ Up to 200 results can be displayed. If the objects you need are not shown in the result list, you can search them by object name.

#### Source Database Object

Search database name, supporting fuzzy match

1 database in total, with 1 displayed More

db-dst

[Refresh](#) [Select all](#) [Clear](#)

#### Selected Object ⓘ

Globally search for original object names, with

☐ db-dst (Entire database selected)

[Unfold all](#) [Fold all](#) [Select all](#) [Clear](#) [Revert to C](#)

Category	Parameter	Description
Data Initialization Option	Initialization Type	<p>In this scenario, select <b>**Structure initialization/Full data initialization**</b>.</p> <p>Structure initialization: Table structures in the source instance will be initialized into the target instance before the sync task runs.</p> <p>Full data initialization: Data in the source instance will be initialized into the target database before the</p>

		sync task runs.
	If Target Already Exists	In this scenario, select <b>**Precheck and report error**</b> . Precheck and report error: If a table with the same name exists in both the source and target databases, an error will be reported, and the task will stop. Ignore and execute: Full and incremental data will be directly added to tables in the target instance.
Data Sync Option	Conflict Resolution Method	Select a conflict resolution policy based on the business conditions. In this scenario, select <b>**Overwrite**</b> . Report: If a primary key conflict is found during data sync, an error will be reported, and the data sync task will be paused. Ignore: If a primary key conflict is found during data sync, the primary key record in the target database will be retained. Overwrite: If a primary key conflict is found during data sync, the primary key record in the source database will overwrite that in the target database.
	SQL Type	Supported operations include INSERT, UPDATE, DELETE, and DDL. If you select <b>**Custom DDL**</b> , you can select different DDL statement sync policies as needed. For more information, see <a href="#">Setting SQL Filter Policy</a> . In two-way sync, you can select <b>**DDL**</b> in at most one task. In this scenario, select <b>**DDL**</b> in task 1 but not task 2.
Sync Object Option	Database and Table Objects of Source Instance	Select the objects to be synced.
	Selected Object	Database/Table mapping (renaming) is supported. Hover over a database or table name, click the displayed <b>**Edit**</b> icon, and enter a new name in the pop-up window.

6. In an A > B forward sync task, DTS will check the source and target database parameters. After all check items are passed, click **Start Task**. In a B > A reverse sync task, DTS will also check the DDL configuration.

**Note:**

If the verification failed, fix the problem as instructed in [Database Connection Check](#) and initiate the verification task again.



If an alarm is displayed in the verification result, it will not affect the task start, but we recommend you click **View Details** to get the suggestions for adjustment.

DDL check

Source and target database parameter check

Create Verification Task		
Query Verification Result		
connect db check	Passed	Passed
necessary check	Passed	Passed
version check	Passed	Passed
source instance privilege check	Alarm	<a href="#">View</a>
simple instance param check	Passed	Passed
target instance privilege check	Passed	Passed
check if target instance has conflict content	Passed	Passed
check if there's enough space in target instance	Passed	Passed
source instance binlog param check	Passed	Passed
foreign key constraint check	Passed	Passed
partial table foreign key constraint check	Passed	Passed
view check	Passed	Passed
warning param check	Alarm	<a href="#">View</a>

7. Return to the data sync task list, and you can see that the task has entered the **Running** status.

## Creating a sync task 2: Reverse sync (B > A)

The operations of forward and reverse sync tasks are basically the same. The following only describes their differences:

1. Confirm the status of task 1. When task 1 enters the "incremental sync" phase, start configuring task 2.

This task configuration timing is required only when database B is empty. In other scenarios, there is no need to wait.

2. Set source and target databases.

Swap the data in source and target databases in task 1.

1

Set source and target databases

>

2

Set sync options and objects

>

3

Verify task

Task Configuration

Task Name \*

sync-

Running Mode \*

Immediate execution

Scheduled execution

Source Instance Settings

Source Instance Type \*

MySQL

Source Instance Region

South China(Guangzhou)

Service Provider \*

Others

AWS

Alibaba Cloud

Access Type \*

Public Network

Self-Build on CVM

Direct Connect

VPN Access

Database

CCN

Cross-/Intra-Account \*

Intra-account

Cross-account

Instance ID \*

cdb-

Account \*

root

Password \*

Please enter password

Test Connectivity

Target Instance Settings

Target Instance Type \*

MySQL

Target Instance Region

South China(Guangzhou)

Access Type \*

Public Network

Self-Build on CVM

Direct Connect

VPN Access

Database

CCN

Instance ID \*

cdb-

Account \*

root

Password \*

Please enter password

Test Connectivity

Database B

Database A

3. Set sync options and objects.

Initialization Type: Do not select.

If Target Already Exists: Select **Ignore and execute**.

Primary Key Conflict Resolution: Select an option based on your business conditions. In this scenario, select **Ignore**.

SQL Type: In two-way sync, you can select DDL in at most one task. In this scenario, select DDL in task 1 but not task 2.

✓ Set source and target databases > 2 Set sync options and objects > 3 Verify task

Data Initialization Option

Initialization Type

Structure Initialization

Full data initialization

If Target Already Exists \*

Precheck and report error

Ignore and execute

Data Sync Option

Primary Key Conflict Resolution \*

Report

Ignore

Overwrite

SQL Type

DML

Insert

Update

Delete

DDL

DDL

Sync Object Option

ⓘ Up to 200 results can be displayed. If the objects you need are not shown in the result list, you can search them by object name.

Source Database Object

Selected Object ⓘ

Search database name, supporting fuzzy match ⓘ

1 database in total, with 1 displayed More

▶ ☒ db-dst

Refresh Select all Clear

Globally search for original object names, with fuzzy match supported ⓘ

☐ db-dst (Entire database selected)

Unfold all Fold all Select all Clear Revert to Original Name

Previous

Save and Go Next

4. On the **Verify task** page, check the DDL configuration.

## Stopping a sync task

If you no longer need a sync task, you can select **More > Stop** in the **Operation** column to stop it.

©2013-2022 Tencent Cloud. All rights reserved.

Page 27 of 60

# Creating Many-to-One Sync Data Structure

Last updated : 2024-07-08 15:54:46

## Overview

Many-to-One sync is to sync the content in multiple source databases to one target database. If you use a single database, you may often need to shard the data due to high load or region issues, but storing the databases/tables of the same type in many databases makes data query inconvenient. The many-to-one sync feature can easily solve this problem.

As a many-to-one sync task consists of multiple one-way sync tasks to establish a many-to-one topology, restrictions on one-way sync and relevant operations must be followed. For more information, see the appropriate sync scenario in [Data Sync](#).

## Notes

During full data sync, DTS consumes certain source database resources, which may increase the load and pressure of the source database. If your database configuration is low, we recommend you sync the data during off-peak hours. To avoid duplicate data, make sure that the tables to be synced have a primary key or non-null unique key. You should plan the data in advance. Each source database is responsible for updating (adding, deleting, and modifying) data with different primary keys so as to avoid problems such as primary key conflict and mutual overwriting of data with the same primary key. If there are duplicate primary keys in multiple source databases for business reasons, select an appropriate conflict resolution method as instructed in [Recommended Configuration for Typical Use Case](#) to make the sync behavior and data meet the expectations.

## Application Restrictions

DDL statements in the configurations of multiple sync tasks should not form a ring.

## DDL Configuration Principles

DDL statements in the configurations of multiple sync tasks should not form a ring; otherwise, they will loop in the system, causing errors.

The same table object in the target database cannot receive DDL sync from multiple source databases; otherwise, such DDL statements may conflict with each other in the target database, causing errors.

In many-to-one sync that combines multiple tables with the same name into one, you can select DDL in only one sync task.

In other types of many-to-one sync tasks (such as a task that combines multiple tables with different names into one database), you can select DDL in each task. In this case, select an appropriate DDL sync policy based on the actual conditions.

During verification, the sync system will judge whether the sync task being created will cause a DDL loop or conflict based on all your other sync tasks and provide prompts for your reference.

## Recommended Configurations for Typical Use Cases

A many-to-one sync task consists of multiple one-way sync tasks to establish a many-to-one topology. The creation steps for each one-way sync task are similar to those for a general one-way sync task. They differ only in the following sync option settings:

The following configurations are recommended for typical use cases for your reference.

Example: a sync task among databases A, B, and C needs to be created, where databases A and B have tables with the same name that need to be synced to database C, task 1 is sync from A to C, and task 2 is sync from B to C. To sync data from more source databases to the target database, simply add sync tasks by referring to task 2.

Scenario	Time Requirements	Sync Task	Initialization Type	If Target Already Exists	Conflict Resolution Method	SQL Type
Scenario 1: databases A and B have database/table structures and data, and database C is empty	Task 2 can be started only after task 1 enters the "incremental sync" phase	Task 1	Structure initialization/full data initialization	Ignore and execute	Select an option as needed. Example: if a primary key conflict occurs, and you want the content of database A to prevail, you need to select <b>**Overwrite**</b> for task 1 and <b>**Ignore**</b> or <b>**Report**</b> for task 2. The conflict resolution	Select DDL in at most one task. For operation types other than DDL, keep them consistent between the other multiple tasks.
		Task 2	Full data initialization	Ignore and execute		
Scenario 2: databases A and B have database/table structures and data, and database C has only database/table structures but no data	None	Task 1	Full data initialization	Ignore and execute		
		Task 2	Same as task 1	Same as task 1		

Scenario 3: databases A, B, and C all have database/table structures and data	None	Task 1	Full data initialization	Ignore and execute	method takes effect only for the data with primary key conflict.	
		Task 2	Same as task 1	Same as task 1		

## Directions

The following uses MySQL two-to-one sync (databases A and B have database/table structures and data, and database C is empty) as an example. The many-to-one sync operations for other databases are similar.

### Creating sync task 1 (database A > database C)

1. Log in to the [data sync purchase page](#), select appropriate configuration items, and click **Buy Now**.
2. After successful purchase, return to the [data sync list](#), and you can see the newly created data sync task. You need to configure it before you can use it.
3. In the data sync list, click **Configure** in the **Operation** column to enter the sync task configuration page.
4. On the sync task configuration page, configure the source and target databases and their accounts and passwords, test the connectivity, and click **Next**.

Category	Parameter	Description
Task Configuration	Task Name	DTS will automatically generate a task name, which is customizable.
	Running Mode	Immediate execution and scheduled execution are supported.
Source Database Settings	Source Database Type	Select the TencentDB instance type selected during purchase, which cannot be changed once configured.
	Source Database Region	Select the TencentDB instance A region selected during purchase, which cannot be changed once configured.
	Service Provider	Others (including TencentDB for MySQL and self-built MySQL), AWS, and Alibaba Cloud are supported.
	Access Type	If <b>Other Cloud Vendors</b> is selected as <b>Service Provider</b> , the access type can be public network; if <b>Others</b> is selected as <b>Service Provider</b> , you need to select an access type according to the database deployment conditions.

		<p>Public Network: self-built database connected through a public IP.</p> <p>Self-Build on CVM: self-built database on CVM.</p> <p>Direct Connect/VPN Access: self-built database connected through a Direct Connect/VPN gateway.</p> <p>VPC: self-built database connected through a VPC.</p> <p>Database: TencentDB database.</p> <p>CCN: self-built database connected through CCN.</p>
Target Database Settings	Target Database Type	Select the target database type, which cannot be changed once configured.
	Target Database Region	Select the target database C region, which cannot be changed once configured.
	Access Type	Select the access type of the target database C.

5. On the **Set sync options and objects** page, set the data initialization, data sync, and sync object options and click **Save and Go Next**.

Category	Parameter	Description
Data Initialization Option	Initialization Type	<p>Structure initialization: table structures in the source database will be initialized into the target database before the sync task runs.</p> <p>Full data initialization: data in the source database will be initialized into the target database before the sync task runs.</p> <p>In this document, select <b>**Structure initialization/Full data initialization**</b>.</p>
	If Target Already Exists	<p>Precheck and report error: if a table with the same name exists in both the source and target databases, an error will be reported, and the task will stop.</p> <p>Ignore and execute: full and incremental data will be directly added to tables in the target database.</p> <p>In this document, select <b>**Ignore and execute**</b>.</p>
Data Sync Option	Conflict Resolution Method	Report: if a primary key conflict is found during data sync, an error will be reported, and the data sync task will be

		<p>paused.</p> <p>Ignore: if a primary key conflict is found during data sync, the primary key record in the target database will be retained.</p> <p>Overwrite: if a primary key conflict is found during data sync, the primary key record in the source database will overwrite that in the target database.</p> <p>Select an option as needed.</p>
	SQL Type	Supported operations: INSERT, UPDATE, DELETE, and DDL. In many-to-one sync, you can select DDL in at most one task. In this document, select DDL in task 1 but not other tasks.
Sync Object Option	Database and Table Objects of Source Database	Select the objects to be synced. You can select databases, tables, and views.
	Selected Object	It displays the selected sync objects, and database/table mapping is supported.

6. On the task verification page, the system will check the DDL configuration first and then check the source and target database parameters. After all check items are passed, click **Start Task**.

**Note:**

If the verification failed, fix the problem as instructed in [Fix for Verification Failure](#) and initiate the verification task again.

If an alarm is displayed in the verification result, it will not affect the task start, but we recommend you click **View Details** to get the suggestions for adjustment.

DDL check

Source and target database parameter check

7. Return to the data sync task list, and you can see that the task has entered the **Running** status.

## Creating sync task 2 (database B > database C)

Configure sync task 2 after the previous sync task enters the **incremental sync** phase.

The operations of tasks 1 and 2 are basically the same. The following only describes their differences:

1. Set the sync source and target databases.

Enter the information of databases A and B in the source and target database settings respectively.

2. Set the sync options and objects.

Initialization Type: select **Full data initialization** only but not **Structure Initialization**.

If Target Already Exists: select **Ignore and execute**.

Conflict Resolution Method: select an option as needed.



SQL Type: do not select DDL. In many-to-one sync, you can select DDL in at most one task. In this document, select DDL in task 1 but not other tasks.

## Stopping sync task

If you no longer need a sync task, you can select **More > Stop** in the **Operation** column to stop it.

# Creating Multi-Site Active-Active IDC Architecture

Last updated : 2024-07-08 15:54:46

## Overview

The multi-site active-active IDC architecture refers to multiple IDCs that are deployed in different regions and provide service concurrently. Data can be synced among them in real time. If a disaster occurs in an IDC, its traffic can be routed to other IDCs to implement quick cross-region failover and guarantee business continuity.

The multi-site active-active IDC architecture is implemented by creating multiple two-way sync tasks, each of which consists of two one-way sync tasks. Therefore, restrictions on one-way sync and relevant operations must be followed. For more information, see the appropriate sync scenario in [Data Sync](#).

## Notes

During full data sync, DTS consumes certain source database resources, which may increase the load and pressure of the source database. If your database configuration is low, we recommend you sync the data during off-peak hours. To avoid duplicate data, make sure that the tables to be synced have a primary key or non-null unique key. You should plan the data in advance. Each IDC is responsible for updating (adding, deleting, and modifying) data with different primary keys so as to avoid problems such as primary key conflict and mutual overwriting of data with the same primary key. If there are duplicate primary keys in multiple source databases for business reasons, select an appropriate conflict resolution method to make the sync behavior and data meet the expectations.

## Application Restrictions

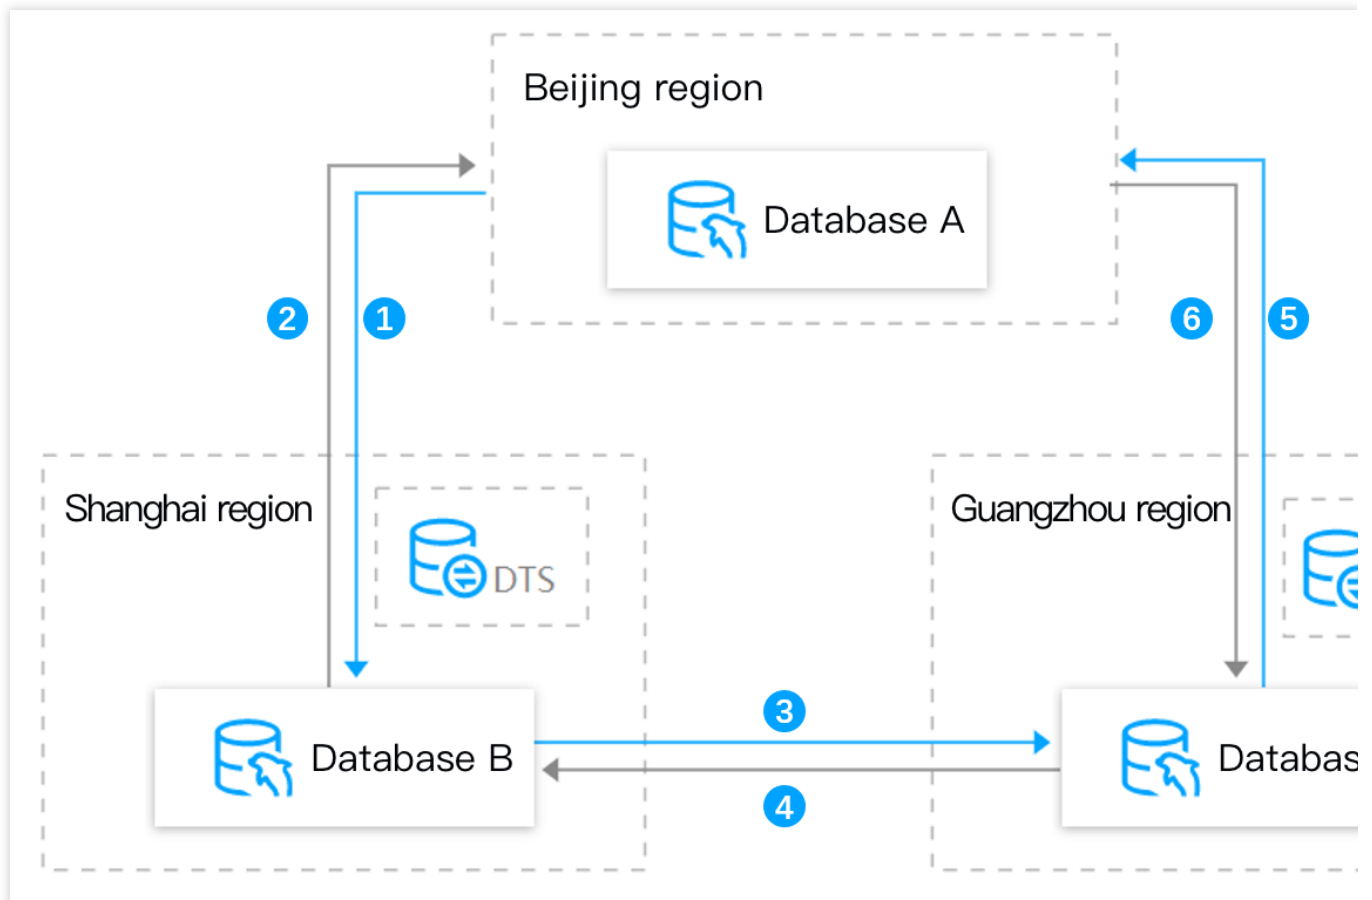
DDL statements in the configurations of multiple sync tasks should not form a ring.

Currently, a two-way sync task can be created between two MySQL databases, two TDSQL-C for MySQL databases, or one MySQL database and one TDSQL-C for MySQL database.

## DDL Configuration Principles

This document uses a specific scenario to describe how to configure DDL statements for easier understanding. For example, in a multi-site active-active IDC architecture, three two-way sync tasks are created among databases

A (Beijing region), B (Shanghai region), and C (Guangzhou region): A <-> B, B <-> C, and C <-> A.



DDL statements in the configurations of multiple sync tasks should not form a ring; otherwise, they will loop in the system, causing errors.

For example, among the three sync tasks (1, 3, and 5) marked by blue lines in the following figure, you can select DDL in up to two of them, and if you select three, a ring will be formed.

The same table object cannot receive DDL sync from multiple IDCs; otherwise, such DDL statements may conflict with each other in the target database, causing errors.

For example, databases A and C have tables with the same name to be synced to database B. Then, you can select DDL in only one task between tasks 1 and 4.

During verification, the sync system will judge whether the sync task being created will cause a DDL loop or conflict based on all your other sync tasks and provide prompts for your reference.

## Recommended Configurations for Typical Use Cases

The multi-site active-active IDC architecture is implemented by creating multiple two-way sync tasks, each of which consists of two one-way sync tasks. Therefore, the operation steps for each sync task in such architecture are basically the same as those for a general one-way sync task. They differ only in the following configurations:

### Sync Option Settings Difference

✓ Set source and target databases

2 Set sync options and objects

3 Verify task

Data Initialization Option

Initialization Type

Structure Initialization

Full data initialization

If Target Already Exists \*

Precheck and report error

Ignore and execute

Data Sync Option

Conflict Resolution Method \*

Report

Ignore

Overwrite

SQL Type

DML \*

Insert

Update

Delete

DDL

DDL

Custom DDL \*

This document recommends the following configuration for a typical multi-site active-active IDC architecture for your reference.

For example, in a multi-site active-active-active IDC architecture, three two-way sync tasks are created among databases A (Beijing region), B (Shanghai region), and C (Guangzhou region): A <-> B (tasks 1 and 2), B <-> C (tasks 3 and 4), and C <-> A (tasks 5 and 6).

Scenario	Time Requirements	Sync Task	Initialization Type	If Target Already Exists	Conflict Resolution Method	SQL Type
Scenario 1: database A has database/table structures and data, and databases B and C are empty	Task 2 can be created only after task 1 enters the "incremental sync" phase	Task 1	Structure initialization/full data initialization	Precheck and report error	Select an option as needed. The conflict resolution method takes effect only for the data with primary key conflict.	Select DDL statements according to the configuration principles. For other operation types, we recommend you keep them consistent among all sync tasks.
		Task 2	Do not select	Ignore and execute		
	Task 4 can be created only after task 3 enters the "incremental sync" phase	Task 3	Structure initialization/full data initialization	Precheck and report error		
		Task 4	Do not select	Ignore and execute		

	Task 6 can be created only after task 5 enters the "incremental sync" phase	Task 5	Structure initialization/full data initialization	Precheck and report error		
		Task 6	Do not select	Ignore and execute		
Scenario 2: databases A, B, and C all have database/table structures and data	None	Tasks 1–6	Full data initialization	Ignore and execute		

## Directions

Creating a multi-site active-active IDC architecture is to create multiple two-way sync tasks. For detailed directions, see [Creating Two-Way Sync Data Structure](#).

# Selecting Data Sync Conflict Resolution Policy

Last updated : 2024-07-08 15:54:46

## Overview

DTS supports complex topology structures, including many-to-one, one-to-many, cascading one-way, two-way, and cascading two-way sync. In such a structure, data is written to multiple nodes at the same time, so primary key conflicts may occur. To address this issue, DTS detects primary key conflicts and provides the following resolution policies:

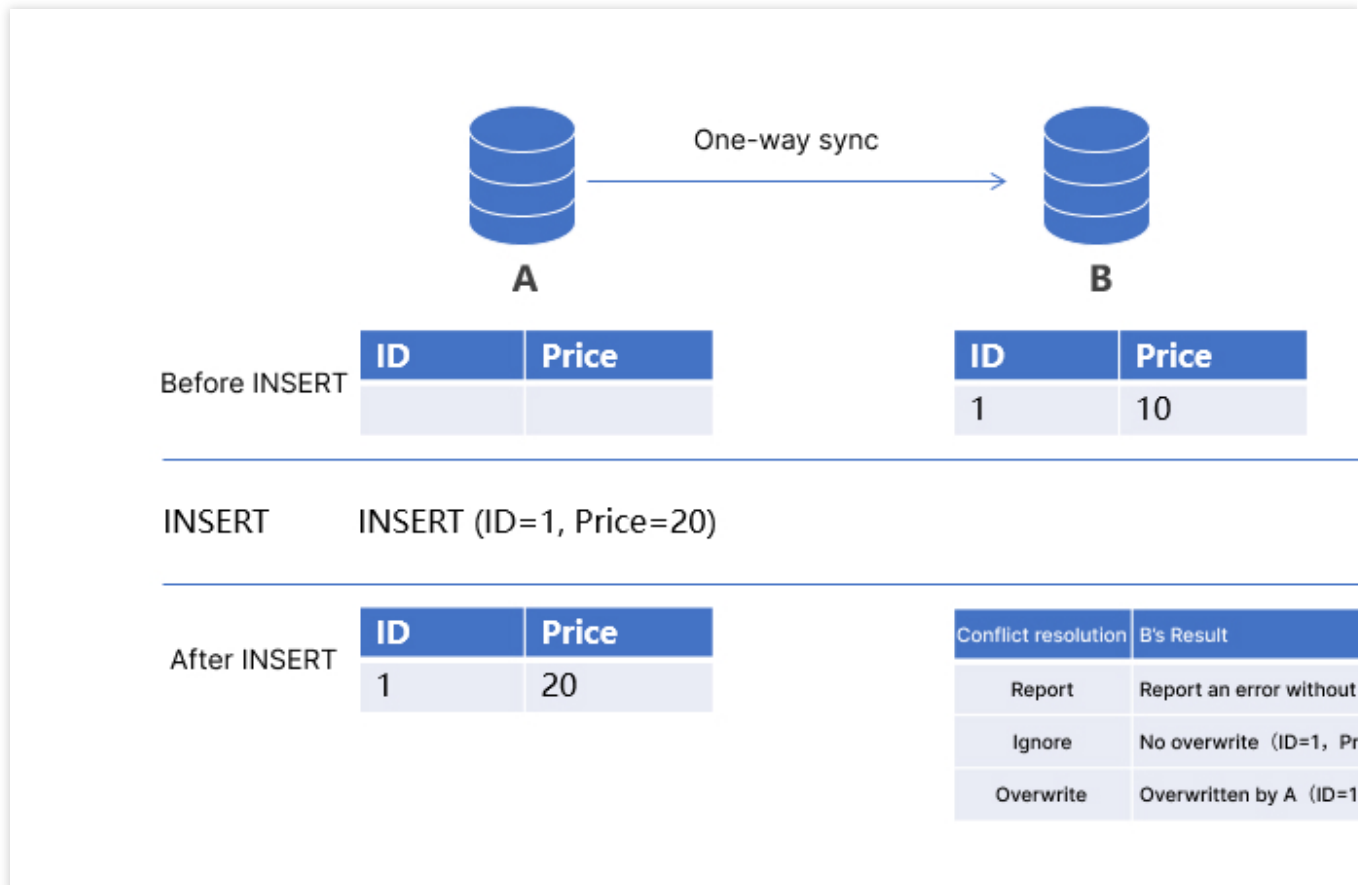
Primary Key Conflict Resolution Policy	Description	SQL Statement Rewrite During Conflict Resolution
Report	During a sync task, if an INSERT statement in the source database has a primary key conflict with the data in the target database, the task will report an error and pause. You need to handle the conflict manually first before proceeding.	The task reports an error, and the SQL statement isn't rewritten.
Ignore	During a sync task, if an INSERT statement in the source database has a primary key conflict with the data in the target database, the data inserted into the source database will be ignored, and the data in the target database will prevail.	If an INSERT statement has a primary key conflict, INSERT will be rewritten to INSERT IGNORE.
Overwrite	During a sync task, if an INSERT or UPDATE statement has a primary key conflict with the data in the target database, the data in the target database will be overwritten by the inserted or updated data in the source database.	If an INSERT or UPDATE statement has a primary key conflict, INSERT or UPDATE will be rewritten to REPLACE INTO or DELETE + REPLACE INTO respectively.

## Examples

Primary key conflict resolution policies take effect only for INSERT and UPDATE primary key conflicts but not in non-conflict scenarios. After a policy is applied, the task can report an error or proceed once a conflict occurs. Below are examples of two primary key conflict scenarios with results under different policies.

## INSERT primary key conflict

An A > B one-way sync with `ID` as the primary key is created. When an INSERT statement in A has a primary key conflict with the data in B during data sync, DTS will handle the conflict according to the configured conflict resolution policy.



The respective sync results in B under different policies are as detailed below:

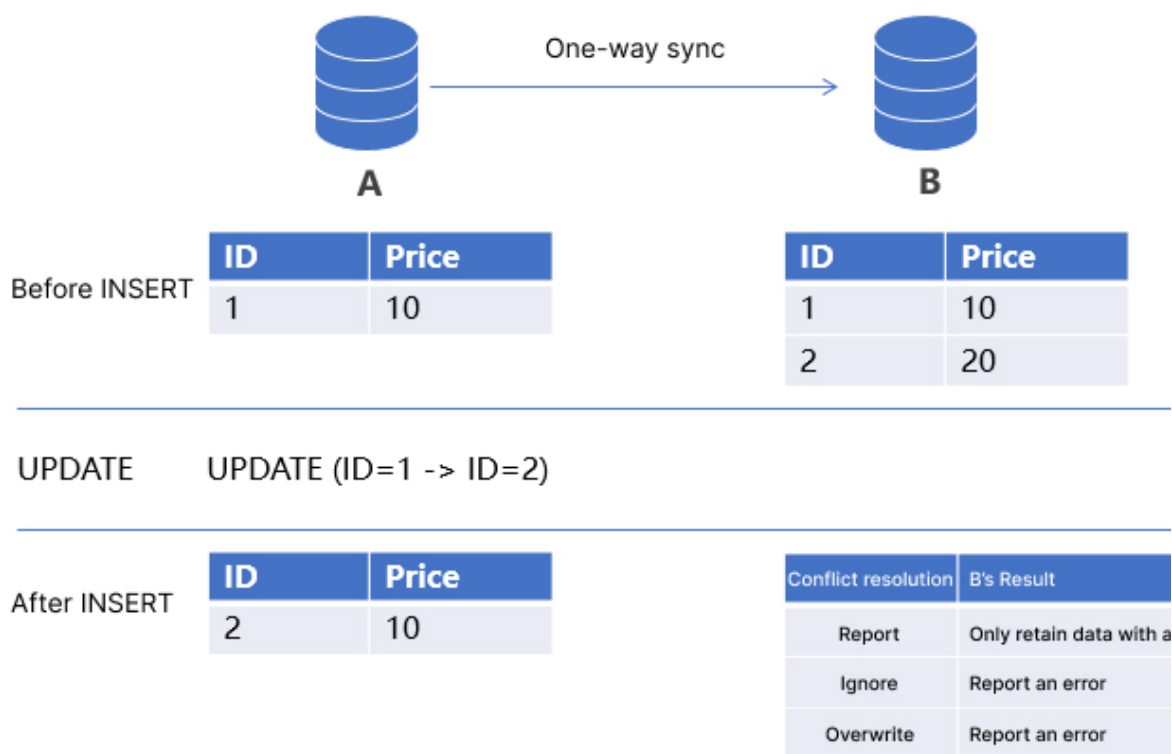
**Report:** The task reports an error, and the data in B remains unchanged (ID=1, Price=10).

**Ignore:** The task ignores the data with the same primary key in A, and the data in B remains unchanged (ID=1, Price=10).

**Overwrite:** The task overwrites the data in B with the data with the same primary key in A, and the data in B becomes ID=1, Price=20 .

## UPDATE primary key conflict

In some scenarios, you may modify the primary key, leading to a primary key conflict. For example, the primary key in A is updated (ID=1 > ID=2), which will conflict with the data with primary key `ID` being `2` in B.



The respective sync results in B under different policies are as detailed below:

Report: The task reports an error, and the data in B remains unchanged.

Ignore: The task reports an error, and the data in B remains unchanged. Note that DTS ignores the conflict in this case.

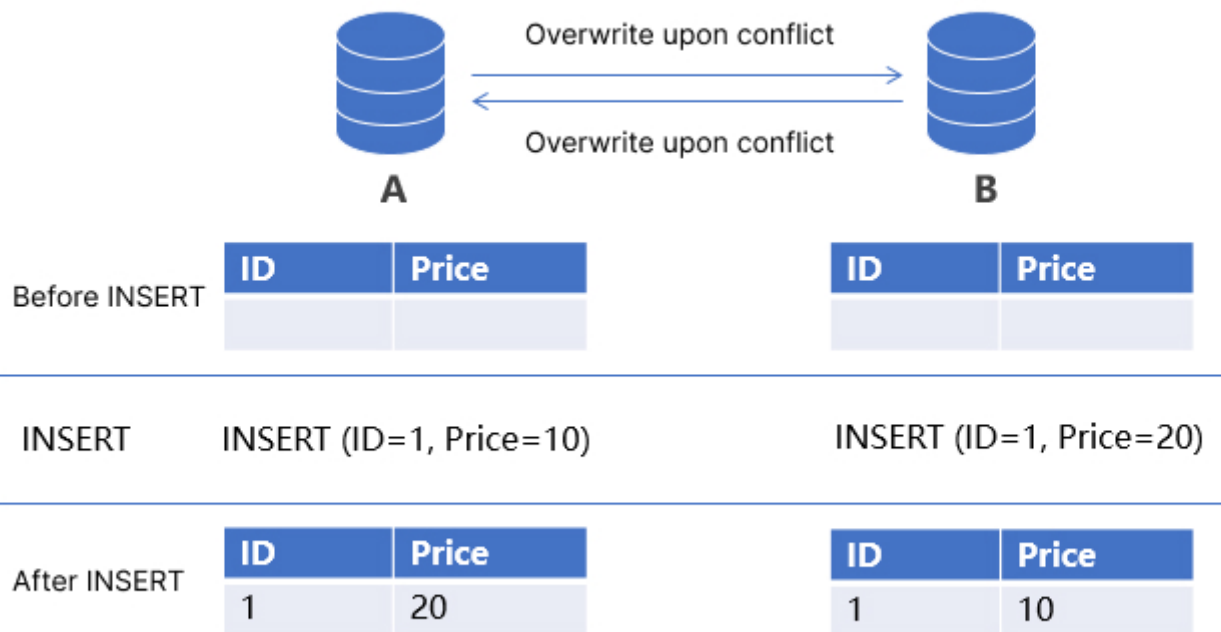
Overwrite: The task overwrites the data in B with the data with the same primary key in A, and only the data with primary key **2** exists in B (ID=2, Price=10).

## Conflict Resolution Policy and Data Consistency

In complex data architectures such as 2-region-3-DC and multi-site active-active architectures, data may need to be written to three or more nodes at the same time, and it is crucial to guarantee the data consistency across multiple nodes. Many users believe that they can use a primary key conflict resolution policy to sync the data on the specified node to other nodes, but this actually doesn't work.

In the following two-way sync scenario, the **Overwrite** policy is set for both  $A > B$  and  $B > A$  sync. If different data records with the primary key **1** are inserted into nodes A and B at the same time, they will be swapped with each other between A and B eventually.





In real-world scenarios, to implement data consistency across nodes, you generally need to partition the database by primary key, introduce additional coordination mechanisms such as data overwriting by version number, and use other methods in addition to a conflict resolution policy.

# Using CLB as Proxy for Cross-Account Database Migration

Last updated : 2024-07-08 15:54:46

## Overview

This document describes how to use CLB as a proxy service to establish a network connection between the source database and DTS. This is suitable for migrating/syncing IDC-based self-built databases or databases in another cloud associated with another Tencent Cloud account to the current account and running various tasks. Below is an example:

VPCs A and B are group company networks, VPC C is a subsidiary network, and account C has no permission to manipulate resources of A and B.

A Direct Connect line is established under account A to connect to the self-built IDC network or third-party cloud vendor network, and account B is connected to VPCs A, B, and C through CCN. Therefore, networks in the dotted box have been interconnected, and account C can access the source database.

Use account C for migration/sync through DTS.

For this scenario, you can associate the source database with a CLB instance. Because CLB can interconnect networks across accounts, you can use the CLB instance as a DTS proxy service for routing and forwarding. Key configuration principles are as follows:

1. Use account C to create a CLB instance.
2. Configure the real server in the CLB instance and bind the source database IP to the real server.
3. Create a migration/sync task and enter the CLB address and port as the IP address and port of the source database.

## Directions

### Creating a CLB instance with account C

1. Log in to the [CLB purchase page](#) with account C.
2. Configure CLB instance parameters and select the **Pay as You Go** billing mode and the **Private Network** type.
3. Return to the **Instance Management** page to view the VIP, which will be used in the subsequent DTS configuration.

<input type="checkbox"/>	ID/Name ↑	Monitor...	Status	VIP	Network...	Network	Health !
<input type="checkbox"/>	lb-test		Normal	1	Public Network		Health c d(Config

## Binding the source database IP to the CLB real server

### Note:

The CLB operations in the following steps are for reference only, subject to the descriptions in [Hybrid Cloud Deployment](#).

1. On the **Instance Management** page in the CLB console, click the ID of the CLB instance just purchased.
2. On the **Basic Info** page, click **Configure** for enabling the feature of binding IPs in another VPC in the **Real Server** section.

← lb-kyqjxnhg

Basic Info

Listener Management

Redirection Configurations

Monitoring

Security Group

Basic Info

Name

ID

Status

VIP

Instance Type

Region

Availability Zone

ISP

Network

Normal

Public Network

Guangzhou

Guangzhou Zone 4

BGP

Access Log

The "Store Logs in COS" feature has been deactivated in all regions. For more information, please see [Deactivation of the "Store Logs in COS" Feature](#).

Cloud Log Service ⓘ Not Enabled

Store Logs in COS ⓘ The "Store Logs in COS" feature will be deactivated in all regions. For more information, please see [Deactivation of the "Store Logs in COS" Feature](#).

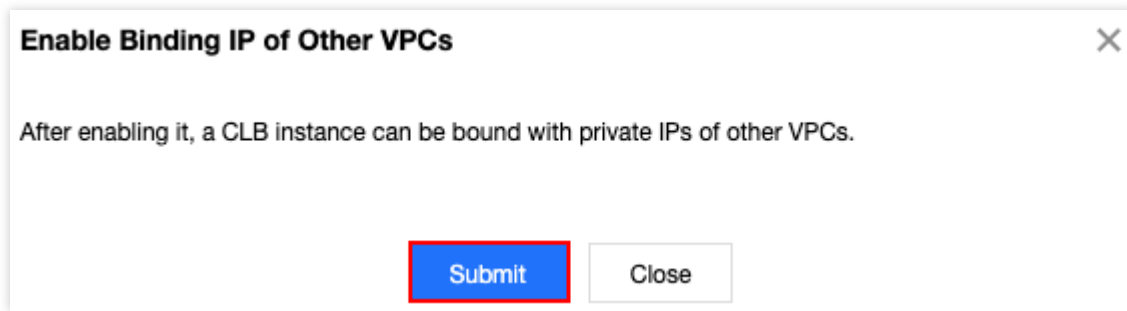
Real Server

Tencent Cloud CLB help you achieve cross-region binding and binding IPs of other VPCs.

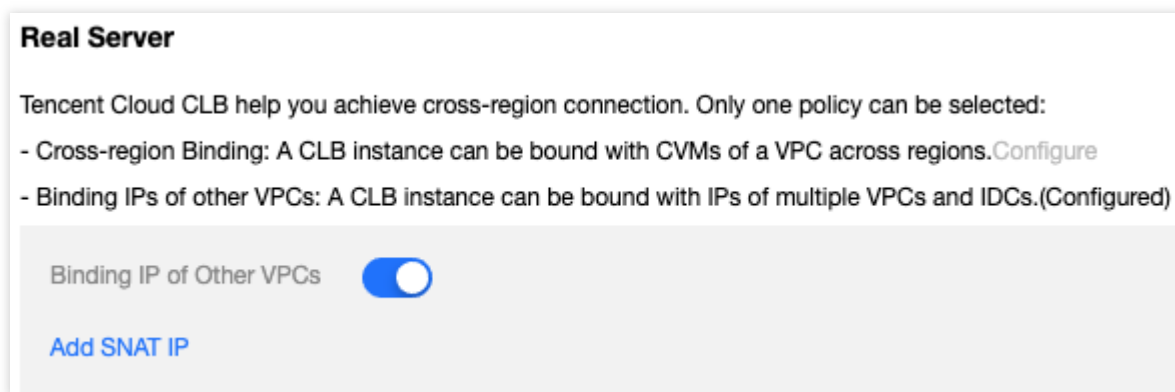
- Cross-region Binding: A CLB instance can be bound to a VPC in another region.

- Binding IPs of other VPCs: A CLB instance can be bound to a VPC in another VPC.

3. In the pop-up window, click **Submit**.



4. After enabling the feature, click **Add SNAT IP** newly displayed in the **Real Server** section.



5. In the pop-up window, select a subnet, click **Add** to assign an IP, and click **Save**.

6. After the SNAT IP is configured.

7. On the instance details page, click the **Listener Management** tab and click **Create** in the **TCP/UDP/TCP SSL/QUIC Listeners** section.

Basic Info **Listener Management** Redirection Configurations Monitoring Security

Note: When custom redirection policies are configured, the original forwarding rules are modified, the redirection policies v

**HTTP/HTTPS Listener**

Create

You've not created any listeners. <a href="#">Create Now</a>	Click to display details
--	--------------------------

**TCP/UDP/TCP SSL Listener**

Create

test1(TCP:10002)	Click to display details
------------------	--------------------------

8. Configure a TCP listener in the pop-up window. You can choose whether to enable health check and session persistence as needed.

9. After configuring the listener, select it and click **Bind** on the right to bind the source database IP address.

Basic Info **Listener Management** Redirection Configurations Monitoring Security Group

Note: When custom redirection policies are configured, the original forwarding rules are modified, the redirection policies will be removed automatically. You need to configure it again. [See details](#)

**HTTP/HTTPS Listener**

Create

You've not created any listeners. [Create Now](#) Click to display details

**TCP/UDP/TCP SSL Listener**

Create

test-tcp(TCP:3306)

**Listener Details** Expand

**Bound Real Server**

Bind Modify Port Modify Weight Unbind

☐ CVM ID/Name Port Health Status<sup>①</sup> IP Address

Listener created. Please [Bound real serv](#)

10. In the pop-up window, select **Another private IP**, enter the source database IP address and port to be bound, set the weight, and click **OK**.

**Bind with backend service**

Target type<sup>①</sup> ☐ Instance ☒ Other Private IP

Default Port  Default Weight

IP	Port	Weight <sup>①</sup>
Source Database IP	1-65535	10

11. Return to the **Real Servers Bound** section to view the bound source database IP.

## Configuring a DTS task

The configuration steps for a DTS task with CLB as an proxy are basically the same as those described in [Migration from MySQL to TencentDB for MySQL](#) or sync from MySQL/MariaDB/Percona to MySQL, with only the following difference:

After purchasing a data migration/sync task with account C, in the **Set source and target databases** step, select **VPC** as the access method (you need to [submit a ticket](#) to enable this option), select the VPC and subnet of account C, and enter the VIP address of the CLB instance as the host address.

1 Set source and target databases >

2 Set migration options and select migration objects >

**Task Configuration**

Task Name \*

Running Mode \* Immediate execution Scheduled execution

**Source Database Settings**

Source Database Type \* MySQL

Service Provider \* Others AWS Alibaba Cloud

Access Type \* Public Network Self-Build on CVM Direct Connect VPN Access Database CCN

Please add the DTS IP addresses to the security group allowlist in advance so that the connectivity test can be quickly

Region South China(Guangzhou)

VPC \* vpc- (Default-VPC) subnet- (Default-Subnet)

Host Address \* Please enter IP address **CLB VIP**

Port \*

Account \*

Password \*

Test Connectivity

# Migrating Self-built Databases to Tencent Cloud Databases via CCN

Last updated : 2024-08-13 14:59:45

## Overview

This document describes how to use the data migration feature of DTS to migrate data from a self-built database to a TencentDB database through CCN.

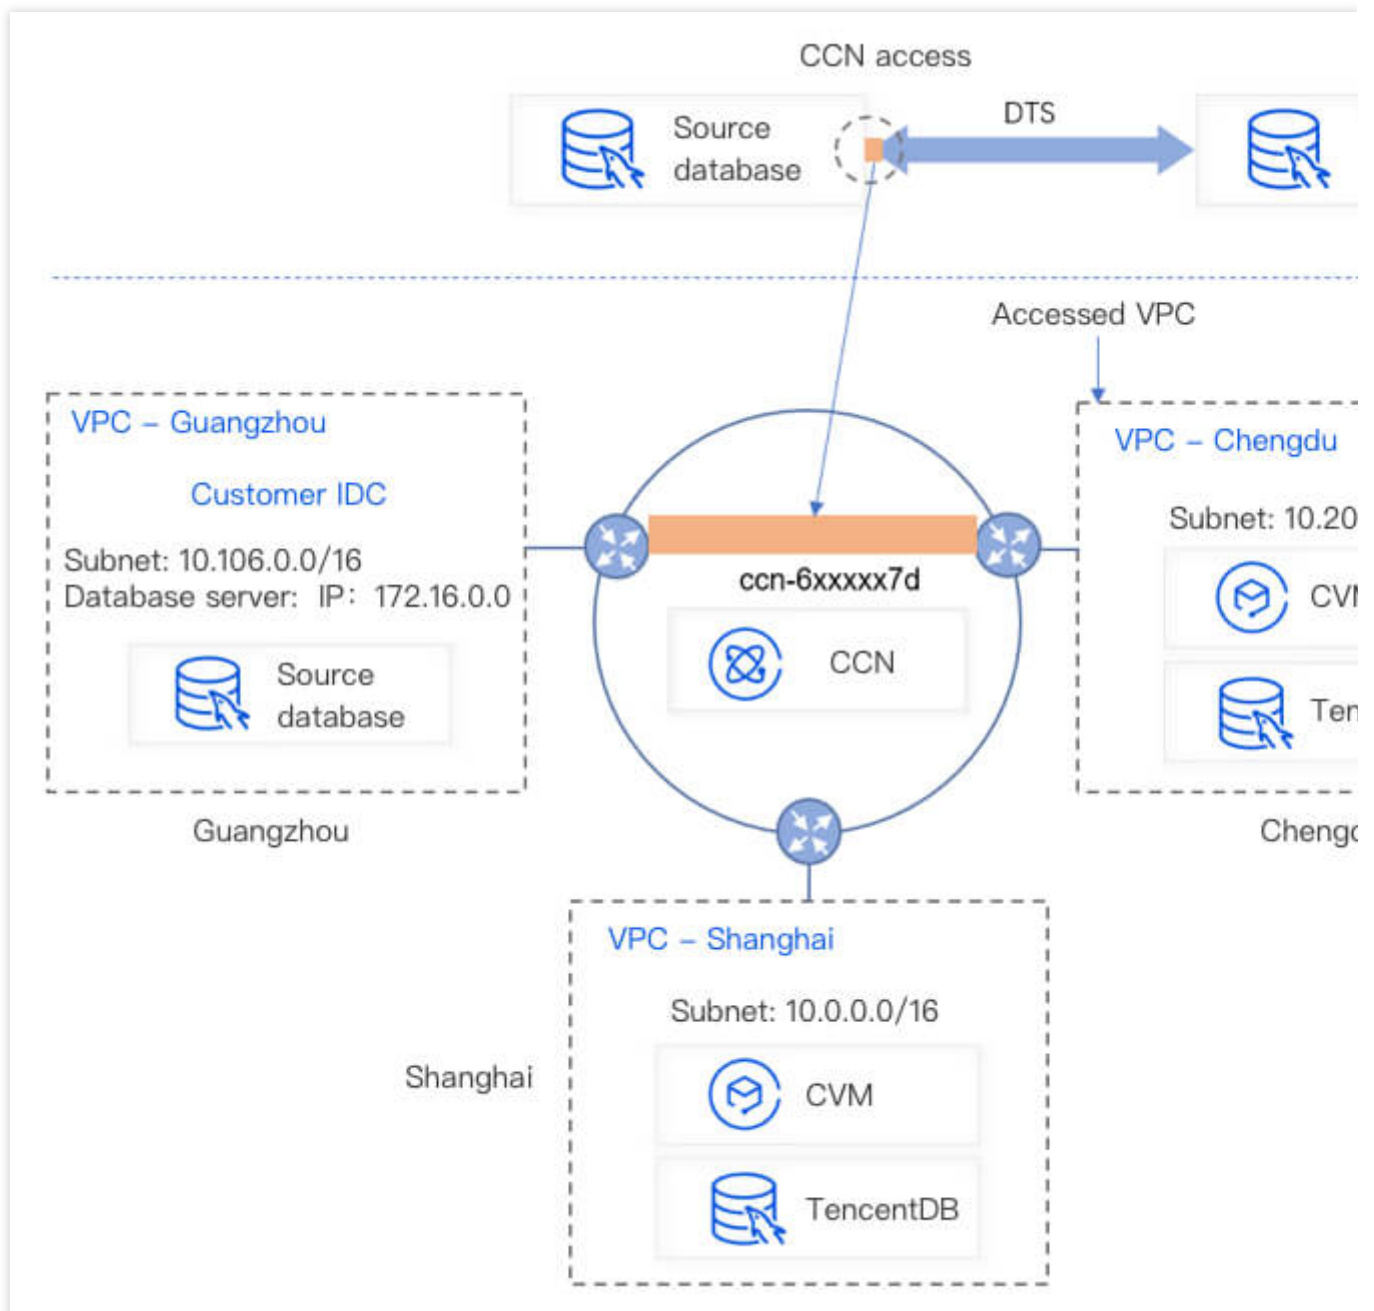
CCN can interconnect a VPC with another VPC or a local IDC. To use CCN access, you must establish cross-VPC and VPC-IDC interconnections through CCN in advance.

In this scenario, you have used CCN to interconnect the three networks of VPC-Guangzhou, VPC-Chengdu, and VPC-Shanghai, have a self-built database in Guangzhou, and plan to migrate the data in the source database in Guangzhou to the target database in Nanjing. VPC-Chengdu is selected as the **Accessed VPC**.

## Configuration Principles

When selecting CCN access, you need to connect the source database to the source of the DTS migration/sync linkage over CCN as follows: source database > accessed VPC > source of the migration/sync linkage, as shown in orange below.





The accessed VPC and the source of the migration/sync linkage are interconnected as follows in the entire DTS task:  
The source of the migration/sync linkage is the network in the region of the source database selected during the task purchase, as shown below:

The region of the source database selected during task purchase must be the same as the region of the accessed VPC; otherwise, the networks cannot be interconnected, and DTS will change the former to the latter.

Creation Mode	<b>Create task</b>	Create similar task
Billing Mode	<b>Pay as you go</b>	
Source Instance Type	<b>MySQL</b>	Redis MongoDB MariaDB PostgreSQL Percona SQL Server
Source Instance Region	<div>South China</div> <div>Guangzhou Shenzhen Finance Shenzhen Shanghai Shanghai Finance Hangzhou Nanjing</div> <div>North China</div> <div>Beijing Tianjin Beijing Finance Singapore Bangkok Jakarta Silicon Valley</div> <div>Europe</div> <div>Frankfurt Moscow Seoul Tokyo Mumbai Virginia</div>	
Target Instance Type	<b>MySQL</b>	TDSQL-C MySQL MariaDB TDSQL MySQL TDSQL TDSore
Target Instance Region	<div>South China</div> <div>Guangzhou Shanghai Hangzhou Nanjing Hong Kong (China) Toronto</div> <div>Southeast Asia</div> <div>Singapore Bangkok Jakarta Silicon Valley Chengdu Chongqing Frankfurt</div> <div>South Asia</div> <div>Mumbai Virginia</div>	
Version	<b>NewDTS</b>	Previous Version
Specification	<b>Xlarge</b>	
Tag	<div>+ Add</div>	
Task Name	<b>Naming after Creation</b>	Name Now
Terms of Service	<input type="checkbox"/> I have read and agree to <a href="#">TENCENT CLOUD TERMS OF SERVICE</a>	
Quantity	<div>- 1 +</div>	

**Accessed VPC:** The accessed VPC refers to the VPC in CCN over which the migration/sync linkage is connected. It can be configured when you set the source and target databases as shown below:

The accessed VPC and the VPC of the source database are interconnected over CCN.

## Notes

When DTS performs full data migration, it will occupy certain source instance resources, which may increase the load of the source instance and the database pressure. If your database has low configurations, we recommend that you migrate data during off-peak hours.

Migration is implemented without locks by default, during which no global lock (the FTWRL lock) is added to the source database, and only tables without a primary key are locked.

## Prerequisites


You have created a TencentDB for MySQL instance as instructed in [Creating MySQL Instance](#).

The source and target databases must meet the requirements for the migration feature and version as instructed in [Databases Supported by Data Migration](#).

You have completed all preparations as instructed in [Overview](#).


The source database must have the following permissions:

Migration of the entire instance:



```
CREATE USER 'migration account'@'%' IDENTIFIED BY 'migration password';  
GRANT RELOAD,LOCK TABLES,REPLICATION CLIENT,REPLICATION SLAVE,SHOW DATABASES,SHOW V  
GRANT ALL PRIVILEGES ON `__tencentdb__`. * TO 'migration account'@'%';  
GRANT SELECT ON *.* TO 'migration account';
```

Migration of specified objects:



```
CREATE USER 'migration account'@'%' IDENTIFIED BY 'migration password';
GRANT RELOAD,LOCK TABLES,REPLICATION CLIENT,REPLICATION SLAVE,SHOW DATABASES,SHOW V
GRANT ALL PRIVILEGES ON `__tencentdb__`.* TO 'migration account'@'%';
GRANT SELECT ON `mysql`.* TO 'migration account'@'%';
GRANT SELECT ON database to be migrated.* TO 'migration account';
```

Permissions required of the target database: ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE USER, CREATE VIEW, DELETE, DROP, EVENT, EXECUTE, INDEX, INSERT, LOCK TABLES, PROCESS, REFERENCES, RELOAD, SELECT, SHOW DATABASES, SHOW VIEW, TRIGGER, and UPDATE.

## Application Restrictions

Basic tables, views, functions, triggers, stored procedures, and events can be migrated, while system tables such as `information_schema` , `sys` , `performance_schema` , `__cdb_recycle_bin__` , `__recycle_bin__` , `__tencentdb__` , and `mysql` cannot.

When views, stored procedures, and functions are migrated, DTS will check whether `user1` corresponding to `DEFINER ( [DEFINER = user1] )` in the source database is the same as the migration account `user2` , and if not, DTS will change the `SQL SECURITY` attribute of `user1` in the target database from `DEFINER` to `INVOKER ( [INVOKER = user1] )` after the migration, and set the `DEFINER` in the target database to the migration account `user2 ( [DEFINER = migration account user2] )`. If the view definition in the source database is too complex, the task may fail.

If the source MySQL database is a non-GTID database, DTS doesn't support HA switch for it. If it is switched, DTS incremental sync may be interrupted.

Only data with the following three database engines can be migrated: InnoDB, MyISAM, and TokuDB. Tables with other engines will be skipped during migration by default.

Correlated data objects must be migrated together; otherwise, migration will fail. Common correlations include table reference by views, view reference by views, and tables correlated through primary/foreign keys.

During incremental migration, if the source database has distributed transactions or generates binlog statements in the `STATEMENT` format, the migration will fail.

In migration without locks (the source database is Alibaba Cloud ApsaraDB RDS for MySQL 5.6, Alibaba Cloud PolarDB for MySQL 5.6, or Amazon RDS for MySQL, and the target database is TencentDB for MySQL), DDL operations are not supported during full migration.

## Operation Limits

During migration, do not perform the following operations; otherwise, the migration task will fail:

Do not modify or delete user information (including username, password, and permissions) in the source and target databases and port numbers.

Do not run distributed transactions in the source database.

Do not write binlog data in the `STATEMENT` format into the source database.

Do not clear binlogs in the source database.

Do not run DDL operations of changing the database/table structure during database/table structure migration or full migration.

Do not delete the system table `__tencentdb__` during incremental migration.

If you only perform full data migration, do not write new data into the source instance during migration; otherwise, the data in the source and target databases will be inconsistent. In scenarios with data writes, to ensure the data

consistency in real time, we recommend you select full + incremental data migration.

## Supported SQL Operations

Operation Type	Supported SQL Operations
DML	INSERT, UPDATE, DELETE, and REPLACE
DDL	TABLE: CREATE TABLE, ALTER TABLE, DROP TABLE, TRUNCATE TABLE, and RENAME TABLE VIEW: CREATE VIEW and DROP VIEW INDEX: CREATE INDEX and DROP INDEX DATABASE: CREATE DATABASE, ALTER DATABASE, and DROP DATABASE

## Environment Requirements

### Note:

The system will automatically check the following environment requirements before starting a migration task and report an error if a requirement is not met. If you can identify the failed check item, fix it as instructed in [Database Connection Check](#); otherwise, wait for the system verification to complete and fix the problem according to the error message.

Type	Environment Requirement
Requirements for source database	<p>The source and target databases can be connected.</p> <p>The server where the source database resides must have enough outbound bandwidth; otherwise, the migration speed will be affected.</p> <p>Requirements for the database parameters:</p> <p>The <code>`server_id`</code> parameter in the source database must be set manually and cannot be 0.</p> <p><code>`row_format`</code> for the source databases/tables cannot be set to <code>`FIXED`</code>.</p> <p>The values of the <code>`lower_case_table_names`</code> variable in both the source and target databases must be the same.</p> <p>The <code>`connect_timeout`</code> variable in the source database must be greater than or equal to 10.</p> <p>We recommend you enable <code>`skip-name-resolve`</code> to reduce the possibility of connection timeout.</p> <p>Requirements for binlog parameters:</p> <p>The <code>`log_bin`</code> variable in the source database must be set to <code>`ON`</code>.</p> <p>The <code>`binlog_format`</code> variable in the source database must be set to <code>`ROW`</code>.</p>

	<p>The <code>`binlog_row_image`</code> variable in the source database must be set to <code>`FULL`</code>. On MySQL 5.6 or later, if the <code>`gtid_mode`</code> variable is not <code>`ON`</code>, an alarm will be triggered. We recommend you enable <code>`gtid_mode`</code>.</p> <p>You cannot set filter conditions with <code>`do_db`</code> and <code>`ignore_db`</code>.</p> <p>If the source database is a slave database, the <code>`log_slave_updates`</code> variable must be set to <code>`ON`</code>.</p> <p>Foreign key dependency:</p> <p>Foreign key dependency can be set to only one of the following three types: <code>`NO ACTION`</code>, <code>`RESTRICT`</code>, and <code>`CASCADE`</code>.</p> <p>During partial table migration, tables with foreign key dependency must be migrated.</p> <p>The migration precision of DTS for data in <code>`FLOAT`</code> type is 38 digits, and for data in <code>`DOUBLE`</code> type is 308 digits. You should check whether this meets your requirements.</p>
Requirements for the target database	<p>The target database version must be equal to or later than the source database version.</p> <p>The size of the target database space must be at least 1.2 times that of the databases/tables to be migrated in the source database. (Full data migration will execute INSERT operations concurrently, causing some tables in the target database to generate data fragments. Therefore, after full migration is completed, the size of the tables in the target database may be larger than that in the source database.)</p> <p>The target database cannot have migration objects such as tables and views with the same name as those in the source database.</p> <p>The <code>`max_allowed_packet`</code> parameter of the target database must be set to 4 MB or above.</p>
Other requirements	<p>The environment variable <code>`innodb_stats_on_metadata`</code> must be set to <code>`OFF`</code>.</p>

## Directions

### Configuring the network interconnection through CCN

Establish interconnections as instructed in [Connecting Network Instances Under the Same Account](#).

#### Note:

CCN only provides bandwidth below 10 Kbps between all regions free of charge. However, DTS requires a higher bandwidth. Therefore, bandwidth configuration in the link is required.

### Configuring a DTS migration task

1. Log in to the [DTS console](#), select **Data Migration** on the left sidebar, and click **Create Migration Task** to enter the **Create Migration Task** page.



2. On the **Create Migration Task** page, select the types, regions, and specifications of the source and target instances and click **Buy Now**.

3. On the **Set source and target databases** page, configure the task, source database, and target database settings. After the source and target databases pass the connectivity test, click **Create**.

**Note:**

If the connectivity test fails, troubleshoot and fix the problem as prompted and as instructed in [Database Connection Check](#) and try again.

Setting Type	Configuration Item	Description
Task Configuration	Task Name	Set a meaningful name for easy task identification.
	Running Mode	Immediate execution: The task will be started immediately after the task verification is passed. Scheduled execution: You need to configure a task execution time and the task will be started automatically then.
	Tag	Tags are used to manage resources by category in different dimensions. If the existing tags do not meet your requirements, go to the console to create more.
Source Database Settings	Source Database Type	The source database type selected during purchase, which cannot be changed.
	Service Provider	Select Others.
	Region	The source database region selected during purchase, which cannot be changed.
	Access Type	Select CCN. For more information on access types, see <a href="#">Overview</a> .
	Host Address	IP address or domain name for accessing the source MySQL database.
	Port	Port for accessing the source MySQL database.
	Account	Account of the source MySQL database, which must have the required permissions.
	Password	Password of the source MySQL database.
	VPC-based CCN Instance	Only VPC-based CCN instance is supported. You need to confirm the network type associated with CCN.
	Accessed VPC	The accessed VPC refers to the VPC in CCN over which the migration/sync linkage is connected. You need to select a CCN-

		<p>associated VPC other than the VPC where the source database resides.</p> <p>To ensure the network connectivity, you must check whether the following key requirements are met:</p> <p>The selected CCN-associated VPC cannot be in the same region as the host address of the source database. If the source database is MySQL in a self-built IDC, this restriction can be ignored.</p> <p>The selected CCN-associated VPC cannot be in the same VPC as the host address of the source database. If the source database is MySQL in a self-built IDC, the selected VPC cannot be the VPC of the Direct Connect gateway associated with the self-built IDC.</p>
	Subnet	<p>Name of the subnet of the selected VPC.If you cannot pull the subnet, there may be a problem with your account. The account of the accessed VPC must be the same as the migration account.For example, to migrate a database under account A to account B, you should use account B to create a task. Therefore, the accessed VPC must be under account B.</p>
	Region of Accessed VPC	<p>The region of the source database selected during task purchase must be the same as the region of the accessed VPC; otherwise, DTS will change the former to the latter.</p>
Target Database Settings	Target Database Type	<p>The target database type selected during purchase, which cannot be changed.</p>
	Region	<p>The target database region selected during purchase, which cannot be changed.</p>
	Access Type	<p>Select Database.</p>
	Database Instance	<p>Select the instance ID of the target TencentDB database.</p>
	Account	<p>The account of the target TencentDB database, which needs to have required permissions.</p>
	Password	<p>Password of the target TencentDB database.</p>

4. On the **Set migration options and select migration objects** page, configure the migration type and objects and click **Save**.

Configuration Item	Description

Migration Type	<p>Select a type based on your scenario.</p> <p>Structural migration: Structured data such as databases and tables in the database will be migrated.</p> <p>Full migration: The entire database will be migrated.</p> <p>Full + incremental migration: The entire database and subsequent incremental data will be migrated. If there are data writes during migration, and you want to smoothly migrate the data in a non-stop manner, select this option.</p>
Migration Object	<p>Entire instance: Migrate the entire database instance excluding the system databases such as `information_schema`, `mysql`, `performance_schema`, and `sys`.</p> <p>Specified objects: Migrate specified objects.</p>
Specified objects	<p>Select the objects to be migrated in Source Database Object and move them to the Selected Object box.</p>

5. On the task verification page, verify the task. After the verification is passed, click **Start Task**.

If the verification failed, fix the problem as instructed in [Database Connection Check](#) and initiate the verification task again.

Failed: It indicates that a check item fails and the task is blocked. You need to fix the problem and run the verification task again.

Alarm: It indicates that a check item doesn't completely meet the requirements, and the task can be continued, but the business will be affected. You need to assess whether to ignore the alarm or fix the problem and continue the task based on the alarm message.

6. Return to the data migration task list, and you can see that the task has entered the **Creating** status. After 1–2 minutes, the data migration task will be started.

Select **Structural migration** or **Full migration**: once completed, the task will be stopped automatically.

Select **Full + Incremental migration**: After full migration is completed, the migration task will automatically enter the incremental data sync stage, which will not stop automatically. You need to click **Complete** to manually stop the incremental data sync.

Manually complete incremental data sync and business switchover at appropriate time.

Observe whether the migration task is in the incremental sync stage and is not in the lag status. If so, stop writing data to the source database for a few minutes.

Manually complete incremental sync when the data gap between the target and the source databases is 0 MB and the time lag between them is 0 seconds.

7. (Optional) If you want to view, delete, or perform other operations on a task, click the task and select the target operation in the **Operation** column. For more information, see [Viewing Task](#).

## Business cutover

---

After the migration task status becomes **Task successful**, you can formally cut over the business. For more information, see [Cutover Description](#).