

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

Best Practices for DTS Performance Tuning

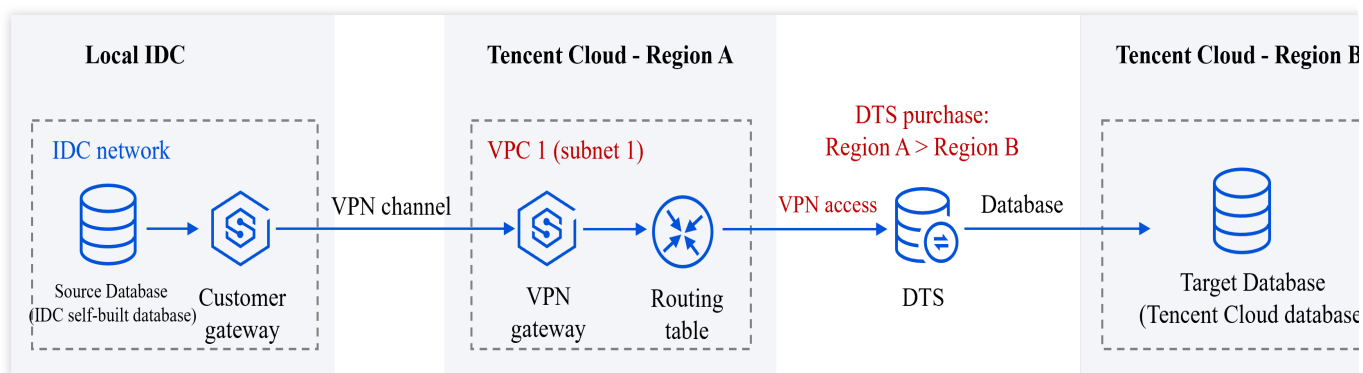
Practical Tutorial

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-built 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.

- 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.
- 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;
```

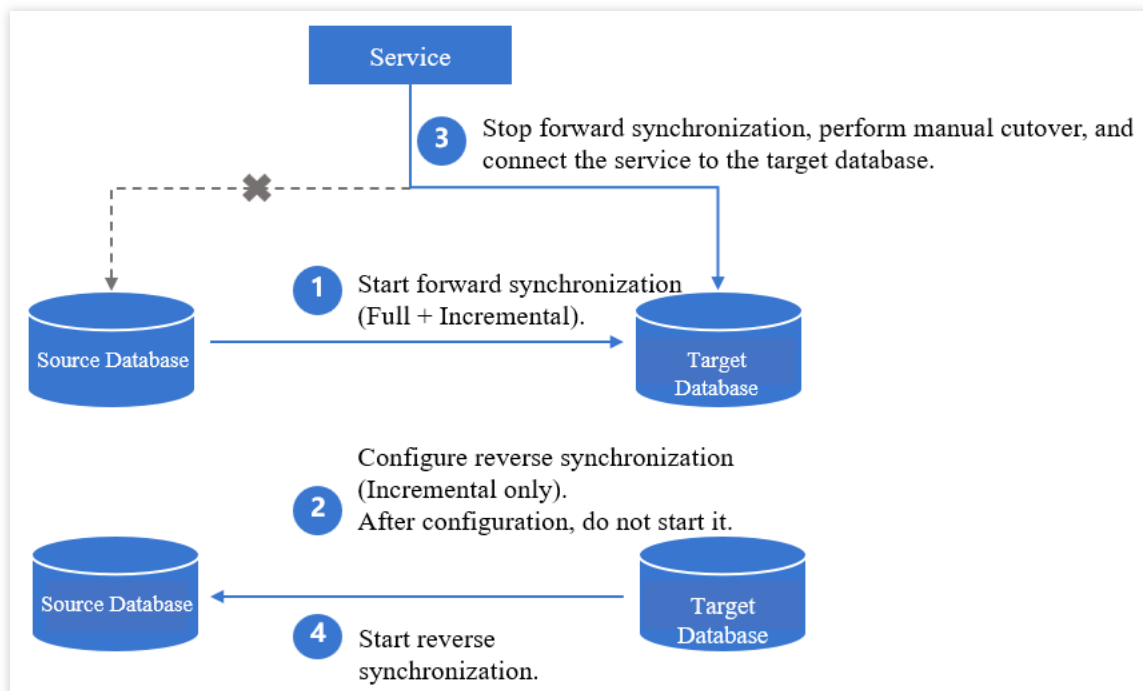
Notes

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 *

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 ✔ 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 *

Account *

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 Others . For a third-party cloud database, select the corresponding service provider. In this scenario, select Others .
	Access Type	Select a type based on your scenario. In this scenario, select Direct Connect or VPN Access , and you need to configure VPN-IDC interconnection as instructed in Direct Connect or VPN Access: Configuring VPN-IDC Interconnection . For the preparations for different access types, see Overview . Public Network: The source database can be accessed through a public IP. Self-Build on CVM: The source database is deployed in a CVM instance. Direct Connect: The source database can be interconnected with VPCs through Direct Connect . VPN Access: The source database can be interconnected with VPCs through VPN Connections . Database: The source database is a TencentDB instance. CCN: The source database can be interconnected with VPCs through CCN . VPC: The source and target databases are both deployed in Tencent Cloud VPCs . To use the VPC access type, submit a ticket 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 Database .
	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: 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 supported

- db-dst (Entire database selected)

[Batch Rename](#)

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

Category	Parameter	Description
Data Initialization Option	Initialization Type	Structure initialization: Table structures in the source instance will be initialized into the target instance before the sync task runs. 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.

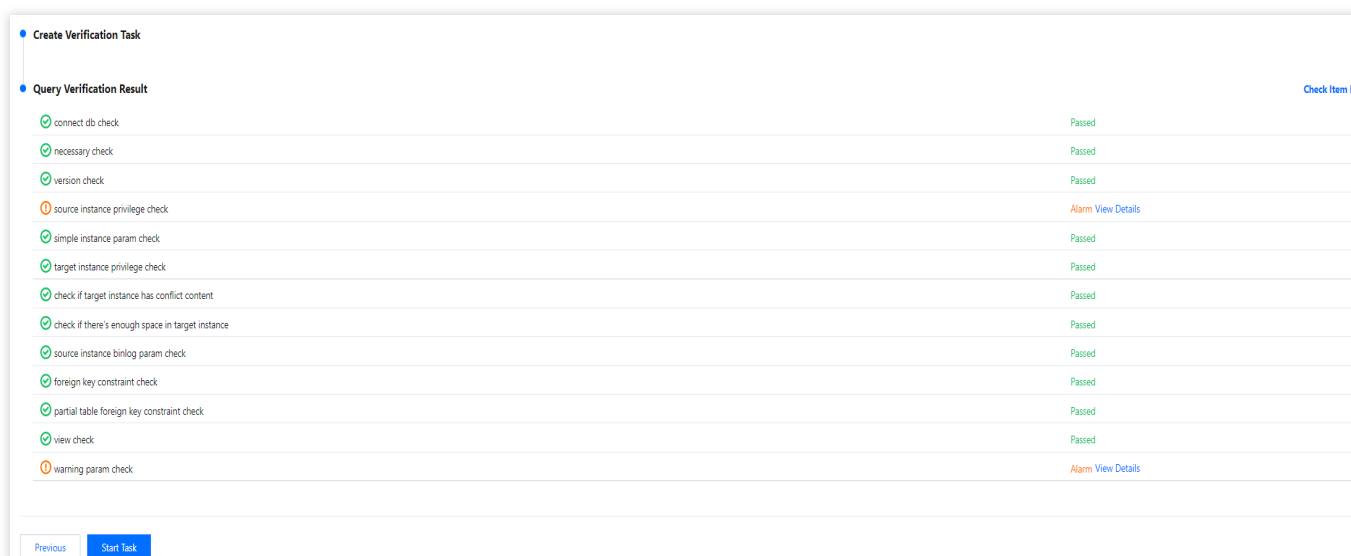
	If Target Already Exists	Precheck and report error: If a table with the same name exists in both the source and target databases, an error will be reported, and the task will stop. Ignore and execute: Full and incremental data will be directly added to tables in the target instance.
Data Sync Option	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 Setting SQL Filter Policy .
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 Syncing Advanced Object .
	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 Syncing Online DDL Temp Table .

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.



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.



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:

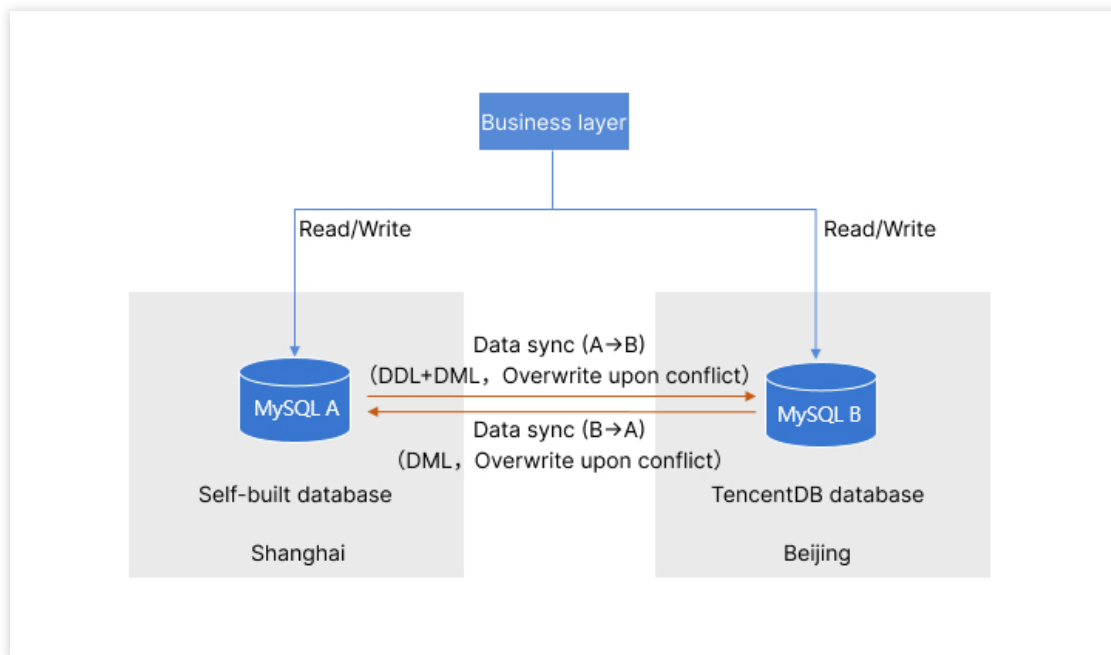
Sync Option Settings Difference

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 **Overwrite** for task 1 and **Ignore** or **Report** 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 *

Database A

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 *

Database B

Test Connectivity

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 Others .
	Access Type	For a third-party cloud database, you can select Public Network generally or select VPN Access , Direct Connect , or CCN based on your actual network conditions. In this scenario, Public Network is selected as an example. For the preparations for different access types, see Overview .
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 Database .
	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

i 1 database in total, with 1 displayed More

- ▶
☑
db-dst

Refresh Select all Clear

Selected Object ⓘ Batch Rename

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

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

		sync task runs.
	If Target Already Exists	In this scenario, select Precheck and report error . 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 Overwrite . 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 Setting SQL Filter Policy . 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.
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 Edit 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.

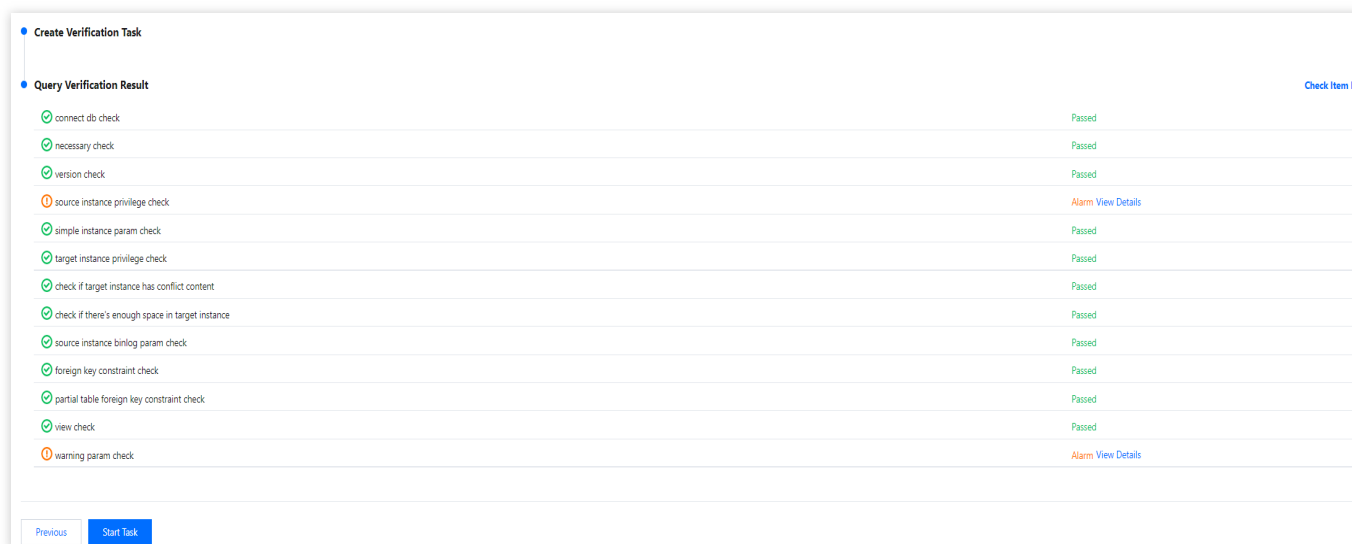
Note:

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



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 *

Running Mode *

Source Instance Settings

Source Instance Type *

Source Instance Region

Service Provider *

Access Type *

Cross-/Intra-Account *

Instance ID *

Account *

Password *

Database B

Target Instance Settings

Target Instance Type *

Target Instance Region

Access Type *

Instance ID *

Account *

Password *

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.

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

Data Initialization Option

Initialization Type:

If Target Already Exists *:

Data Sync Option

Primary Key Conflict Resolution *:

SQL Type

DML:

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

Search database name, supporting fuzzy match

ⓘ 1 database in total, with 1 displayed More

- db-dst

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

Selected Object ⓘ

Batch Rename ⓘ

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](#)

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.

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 **Overwrite** for task 1 and **Ignore** or **Report** 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 Other Cloud Vendors is selected as Service Provider , the access type can be public network; if Others is selected as Service Provider , 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 **Structure initialization/Full data initialization**.</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 **Ignore and execute**.</p>
Data Sync Option	Conflict Resolution Method	<p>Report: if a primary key conflict is found during data sync, an error will be reported, and the data sync task will be</p>

		<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

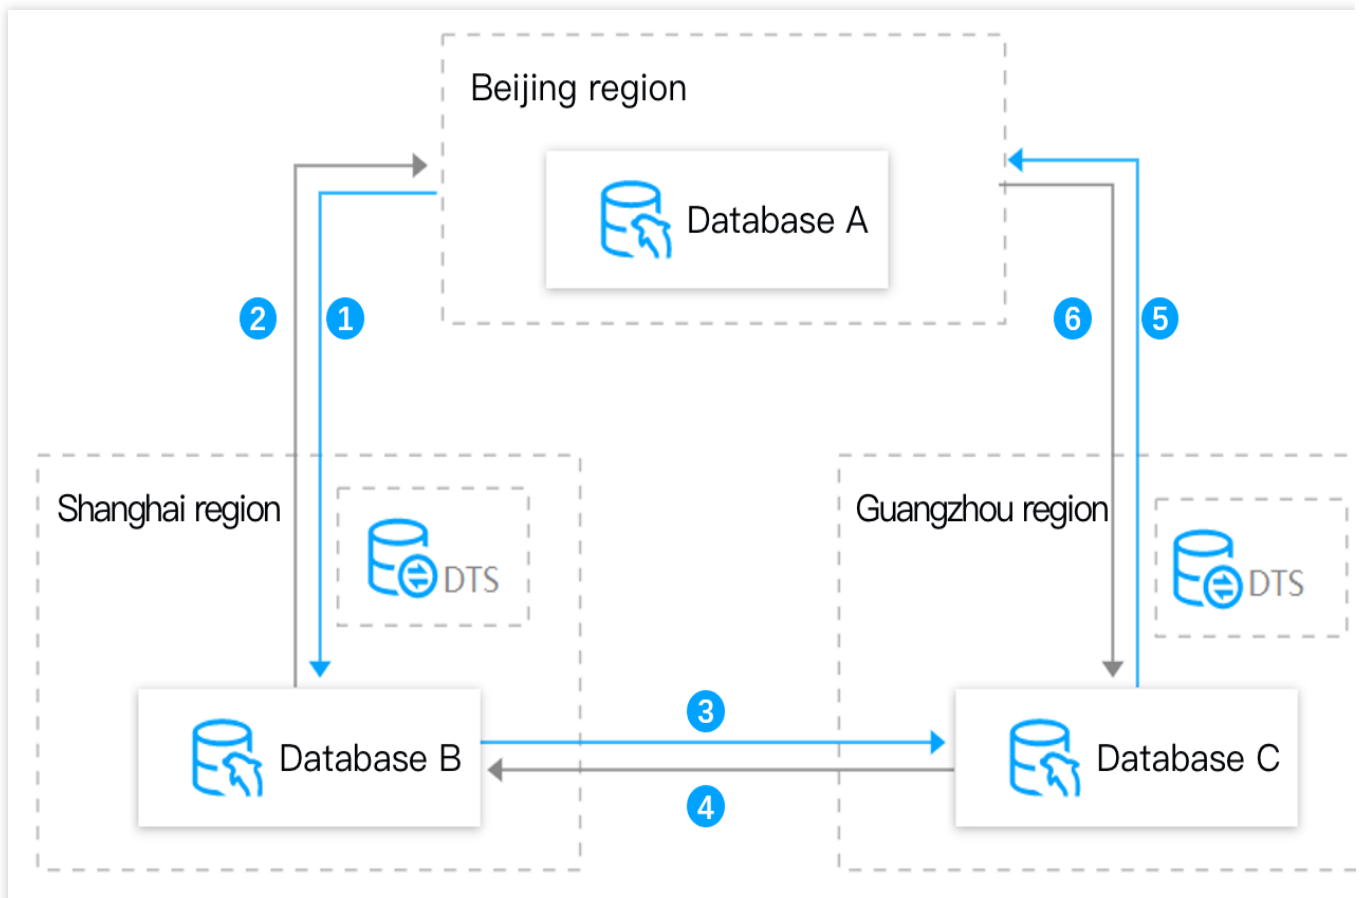
DDL statements in the configurations of multiple sync tasks should not form a ring.

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-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

The screenshot shows the configuration interface for a sync task. It is divided into three steps: 1. Set source and target databases, 2. Set sync options and objects (current step), and 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:

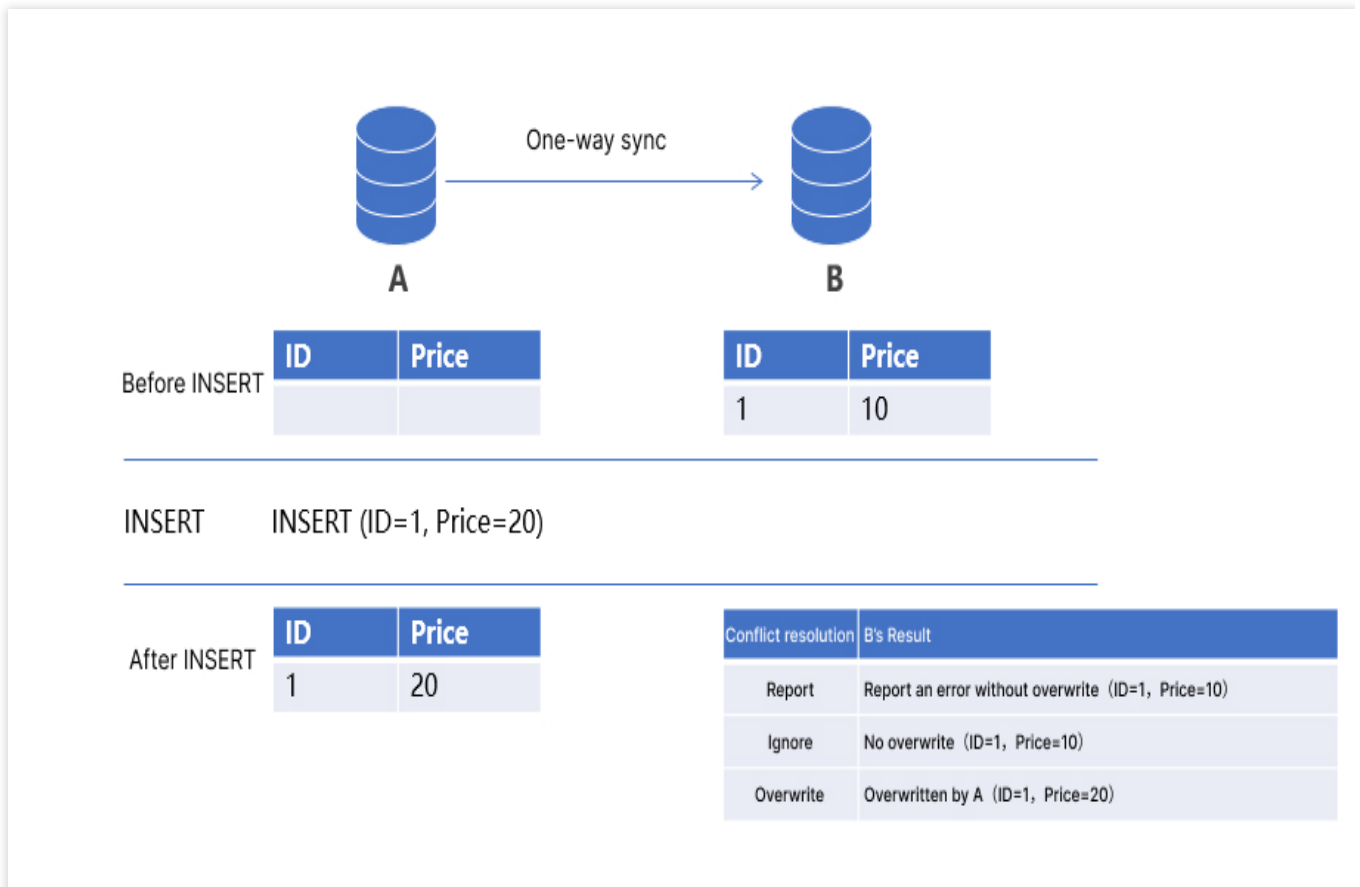
Primary Key Conflict Resolution Policy	Description	SQL Statement Rewrite During Conflict Resolution
Report	During a sync task, if an INSERT statement in the source database has a primary key conflict with the data in the target database, the task will report an error and pause. You need to handle the 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:

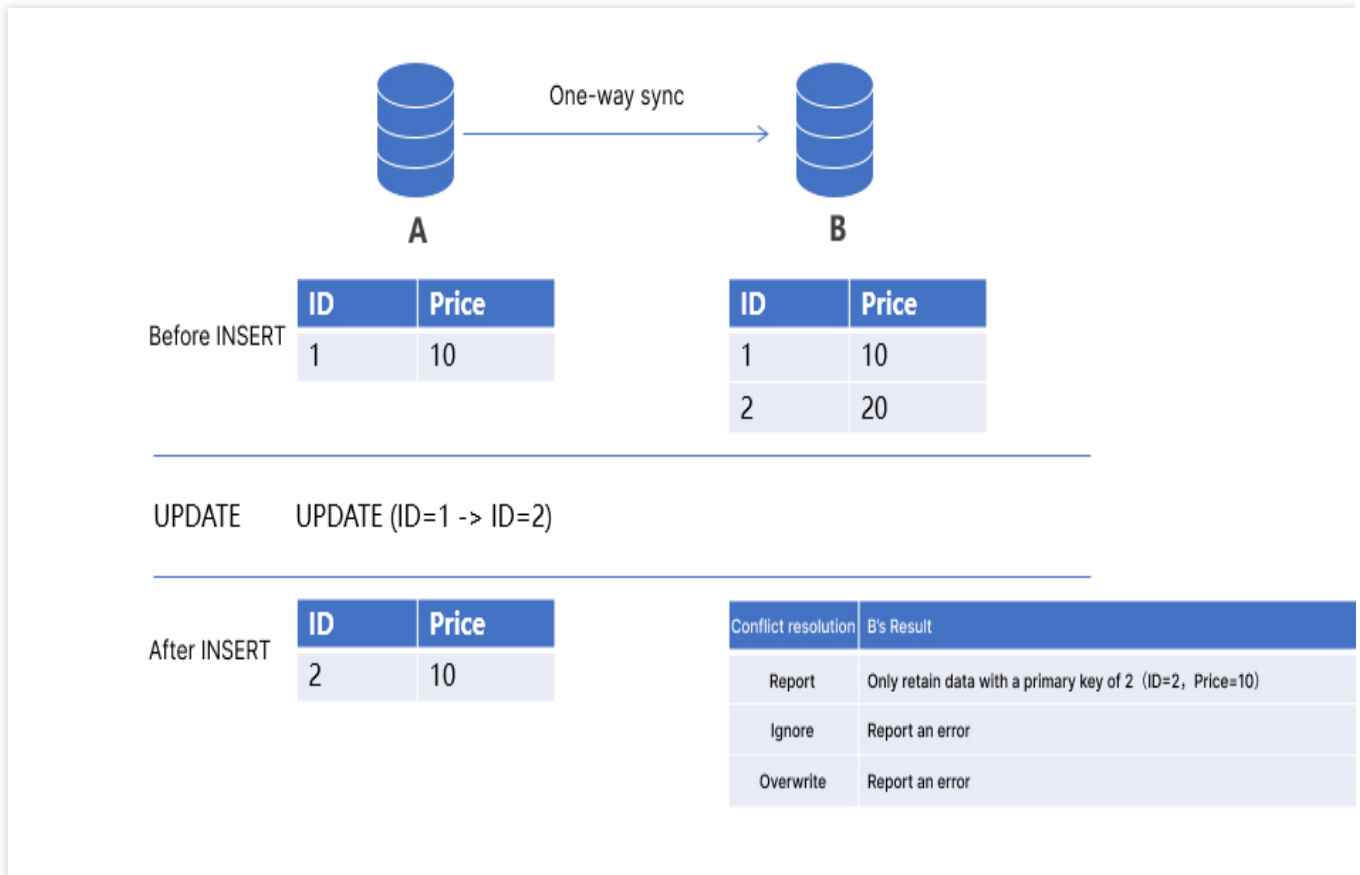
Report: The task reports an error, and the data in B remains unchanged (ID=1, Price=10).

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.

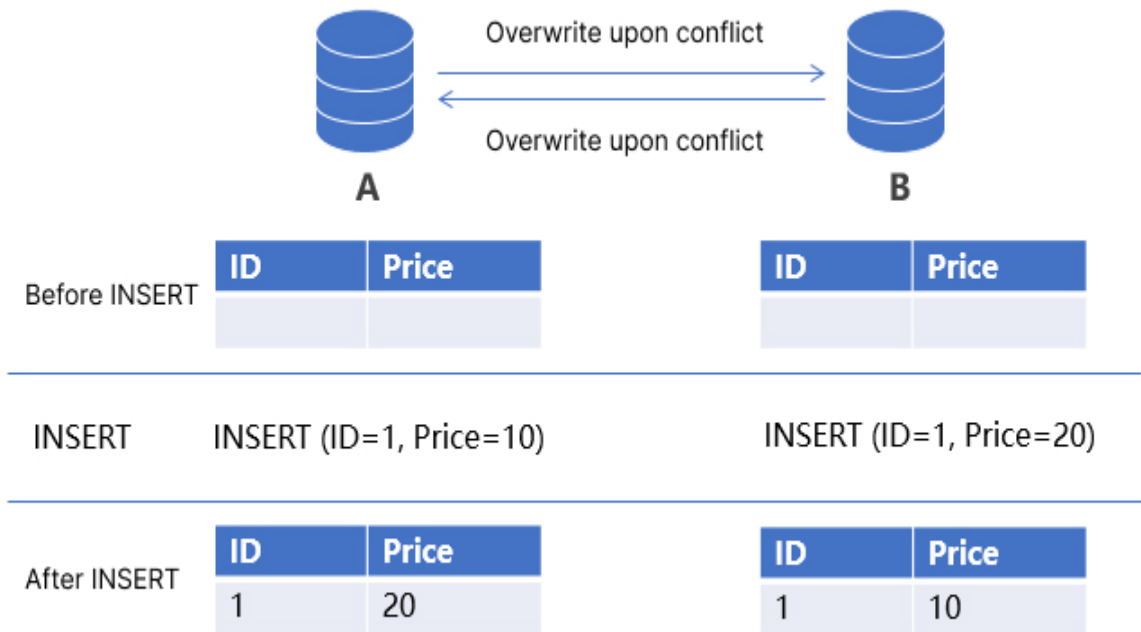
Ignore: The task reports an error, and the data in B remains unchanged. Note that DTS ignores the conflict in this case.

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 : 2024-12-25 14:58:30

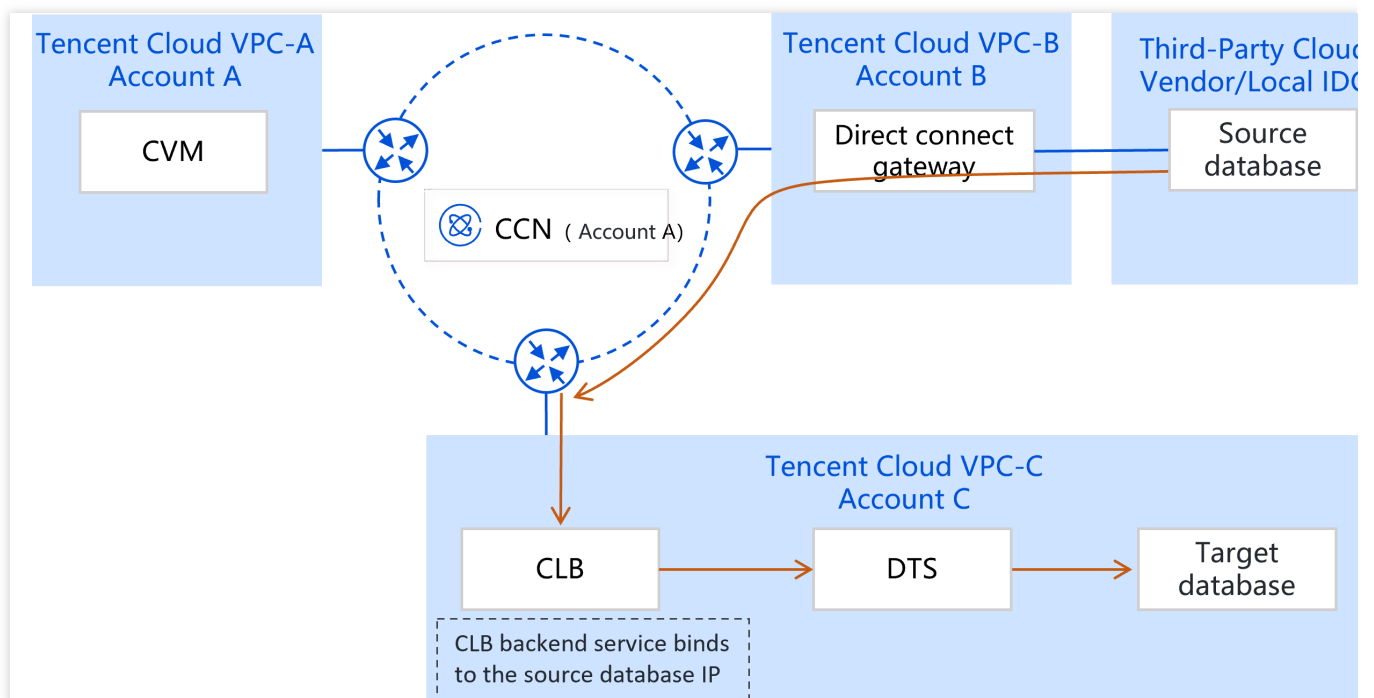
Overview

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. (If you need to migrate/synchronize cloud database instances from another Tencent Cloud account to the current account, there is no need to follow the complex operations in this section. Instead, see [Cross-Account Synchronization of Cloud Database Instances](#) for guidance.)

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

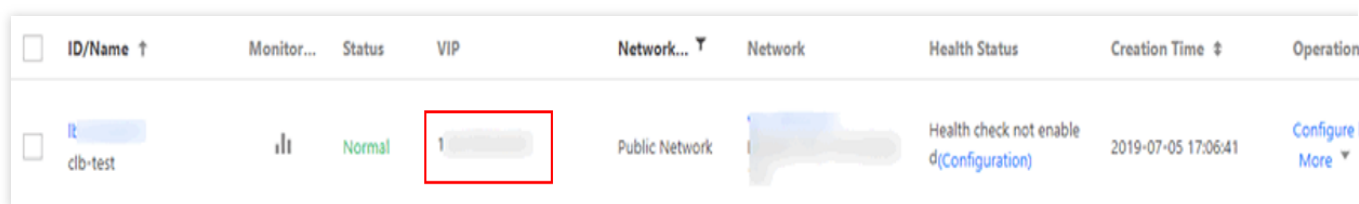
configuration principles are as follows:

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.



ID/Name ↑	Monitor...	Status	VIP	Network... ▾	Network	Health Status	Creation Time ↕	Operation
<input type="checkbox"/> clb-test		Normal	1	Public Network		Health check not enable d(Configuration)	2019-07-05 17:06:41	Configure More ▾

Binding the source database IP to the CLB real server

Note:

The CLB operations in the following guidance are for reference only. If there are differences in the actual console interface, see [CLB official documentation](#) to submit a request. Similarly, any provided screenshots are for reference only. If there are discrepancies, follow the actual console interface.

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: [Redacted]

ID: [Redacted]

Status: Normal

VIP: [Redacted]

Instance Type: Public Network

Region: Guangzhou

Availability Zone: Guangzhou Zone 4

ISP: BGP

Network: [Redacted]

Access Log

The "Store Logs in COS" feature has been unavailable for this region since April 26, 2020. The feature will be deactivated in all regions from June 30, 2020. Please use Cloud Log Service. For more information, please see [Deactivation of CLB "Store Logs in COS" Feature](#)

Cloud Log Service ⓘ Not Enabled ✎

Store Logs in COS ⓘ The "Store Logs in COS" feature is not available in the current region. You can store your access logs to [Cloud Log Service](#)

Real Server

Tencent Cloud CLB help you achieve cross-region connection. Only one policy can be selected:

- Cross-region Binding: A CLB instance can be bound with CVMs of a VPC across regions. [Configure](#)
- Binding IPs of other VPCs: A CLB instance can be bound with IPs of multiple VPCs and IDCs. [Configure](#)

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

Enable Binding IP of Other VPCs ✕

After enabling it, a CLB instance can be bound with private IPs of other VPCs.

Submit **Close**

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

Real Server

Tencent Cloud CLB help you achieve cross-region connection. Only one policy can be selected:

- Cross-region Binding: A CLB instance can be bound with CVMs of a VPC across regions. [Configure](#)
- Binding IPs of other VPCs: A CLB instance can be bound with IPs of multiple VPCs and IDCs. [\(Configured\)](#)

Binding IP of Other VPCs

[Add SNAT IP](#)

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 Group

Note: When custom redirection policies are configured, the original forwarding rules are modified, the redirection policies will be removed automatically.

HTTP/HTTPS Listener

Create

You've not created any listeners. [Create Now](#) [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.

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**.

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 > 3 Verify task

Task Configuration

Task Name *

Running Mode *

Source Database Settings

Source Database Type *

Service Provider *

Access Type * [Access Type Description](#)

Please add the DTS IP addresses to the security group allowlist in advance so that the connectivity test can be quickly passed. For details, see [here](#).

Region

VPC * 4077 available IPs

Host Address * **CLB VIP**

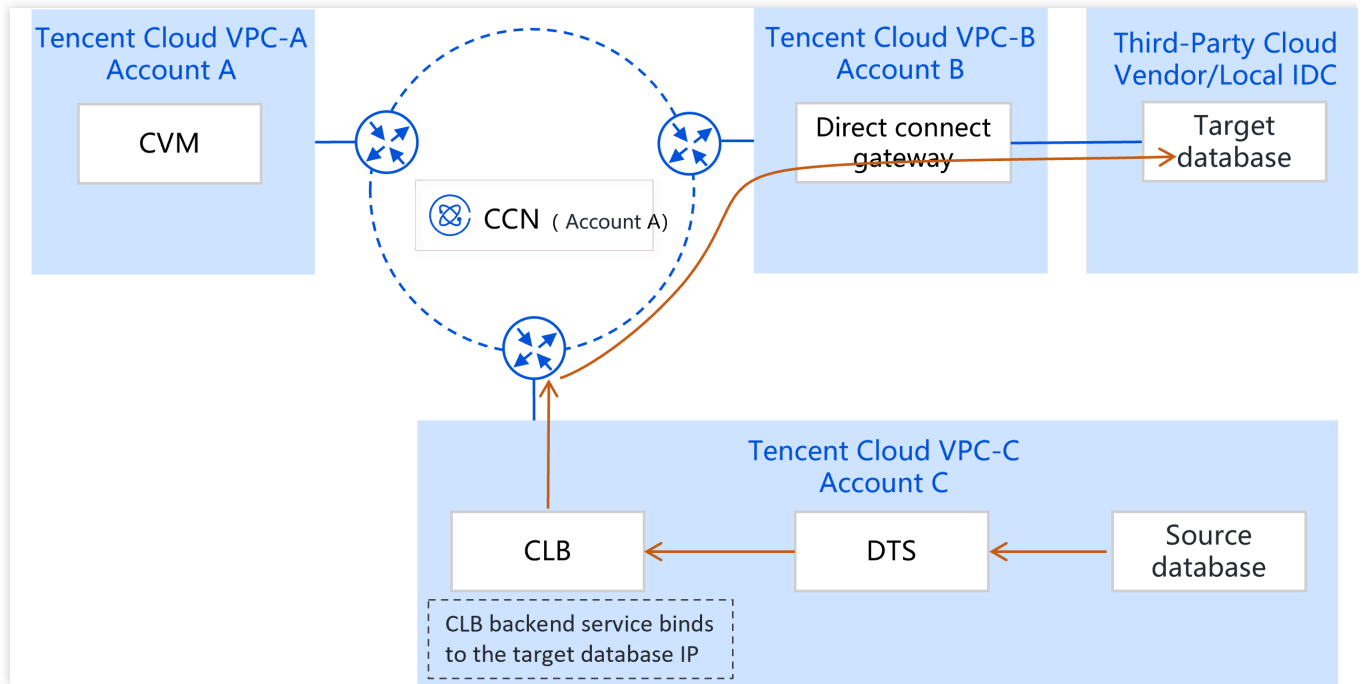
Port *

Account *

Password *

Similar Scenario

Compared with the above scenario, the following scenario is in the opposite migration/sync direction, that is, migrating/syncing databases under the current account to IDC-based self-built databases under other Tencent Cloud accounts or databases provided by third-party cloud vendors. Below are the differences.



1. Create a CLB instance with account C.
2. Configure the real server in the CLB instance and bind the target database IP to the real server.
3. Create a migration/sync task.

Select **Database** as the source database access method.

Select **VPC** as the target database access method. The host address is the same as the CLB address.

Migrating Self-Built Databases to Tencent Cloud Databases via CCN

Last updated : 2024-09-10 17:49:43

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**.

Prerequisites

1. Prepare the target database.
2. Prepare and authorize the account that will execute the DTS task.

2.1 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 VIEW,PROCESS ON *.* TO 'migration account'@'%';
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 VIEW,PROCESS ON *.* TO 'migration account'@'%';
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';
```

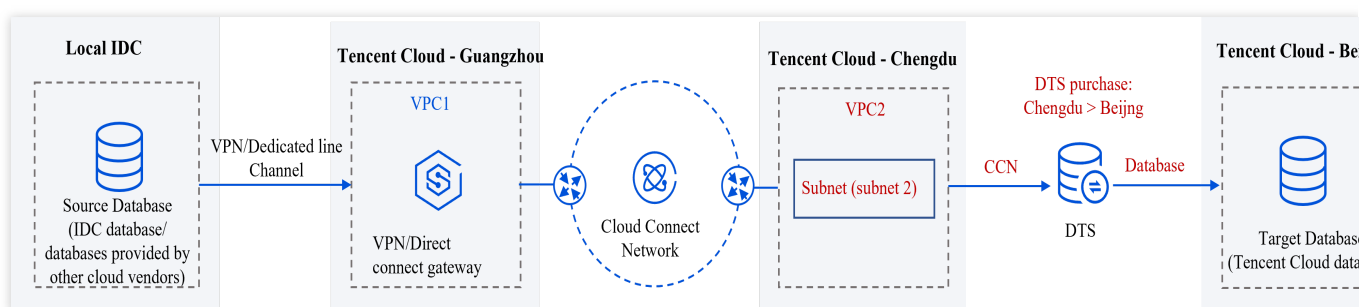
2.2 Permissions required of the target database: ALTER, ALTER ROUTINE, CREATE, CREATE ROUTINE, CREATE TEMPORARY TABLES, CREATE USER, CREATE VIEW, DELETE, DROP, EVENT, EXECUTE, INDEX,

INSERT, LOCK TABLES, PROCESS, REFERENCES, RELOAD, SELECT, SHOW DATABASES, SHOW VIEW, TRIGGER, and UPDATE.

Migration Using the Same-Account CCN

Network Connectivity

Same-Account CCN: This refers to using the CCN under the primary account of the target database to connect with the source database.



1. Use CCN to establish network connectivity.

1.1 Connect the user's source database to VPC1. For detailed instructions, see [Establishing a Connection from VPC to IDC](#).

1.2 Establish connectivity between VPC1 and VPC2. For detailed instructions, see [Establishing interconnectivity between different networks under the same account through CCN](#).

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.

Purchasing the DTS 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**.

For **Source Instance Region**, select the region where the source database's connected VPC (VPC2) is located, which is Chengdu. For **Target Instance Region**, select the region where the target database is located, which is Beijing.

Configuring the DTS Task

1. 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**.

In Source Database Settings, **Access Type** select "CCN", **CCN Instance Account** select "My Account", **CCN Association VPC** select "VPC2" and choose a subnet "subnet2"; In Target Database Settings, **Access Type** select "Database".

CCN Association VPC refers to the VPC within the CCN that is connected to the DTS linkage. You need to select a VPC from all the VPCs connected through the CCN, excluding the VPC that the source database is connected to.

When you select a subnet, if it cannot be pulled, it might be an account issue. The **CCN Association VPC** account and the DTS task account need to be the same. For example, if you want to migrate a database instance from Account A to Account B, the task should be created using Account B. Therefore, the **CCN Association VPC** should be under Account B.

VPC Region: No configuration is required, but the region needs to match the region of the **CCN associated VPC**. If the user did not select the correct region when purchasing the DTS task, DTS will automatically adjust it to ensure the consistency.

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

3. 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.

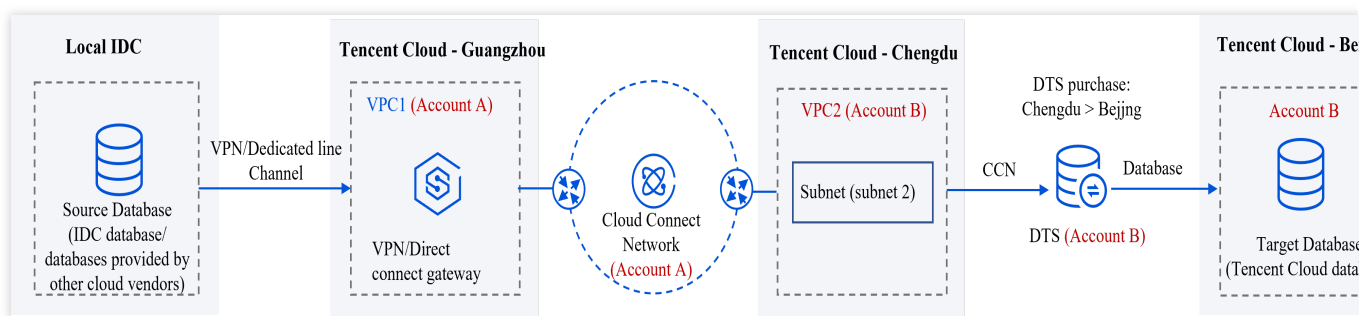
Alarm: It indicates that a check item doesn't completely meet the requirements, and the task can be continued, but the business will be affected. You need to assess whether to ignore the alarm or fix the problem and continue the task based on the alarm message.

4. 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.

Cross-Account CCN Migration Operation

Network Connectivity

Cross-Account CCN: This involves connecting the source database using a CCN under a different Tencent Cloud account.



1. Establish CCN network connectivity.

1.1 Connect the user's source database to VPC1. For detailed instructions, see [Establishing a Connection from VPC to IDC](#).

1.2 Establish connectivity between VPC1 and VPC2. For detailed instructions, see [Establishing interconnectivity between different networks under different accounts through CCN](#).

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.

Authorizing Cross-Account CCN (Optional)

This step is required when using a CCN under a different account. It is not needed for the same-account CCN.

1. Create a DTS role carrier, and select DTS_QCSRole as the role name. For detailed instructions, see [Authorize DTS to access resources of other cloud services](#).

2. [Submit a ticket](#) to apply, after which you will be able to use the CCN feature under a different account for data synchronization in the DTS console.

Purchasing the DTS 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**.

For **Source Instance Region**, select the region where the source database's connected VPC (VPC2) is located, which is Chengdu. For **Target Instance Region**, select the region where the target database is located, which is Beijing.

Configuring the DTS Task

1. After the purchase is completed, the page will automatically redirect to the data migration task list. Select the task you just purchased to start configuration.

2. On the set source and target database page, complete task settings, source database settings, and target database settings. Once the connectivity test between the source and target databases passes, click **Save**.

In the source library settings, select CCN for **Access Type**, select Other Account for **CCN Instance Account Type**, and select VPC2 and a subnet, subnet2 for **CCN Instance Association VPC**. In the target library settings, select Database for **Access Type**.

Note:

Currently, to use a CCN under another account for a DTS task, you need to first [submit a ticket](#) to apply.

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

4. 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.

5. 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.

Best Practices for DTS Performance Tuning

Last updated : 2024-09-11 10:33:30

This document provides guidance on how to improve the DTS transmission rate in various scenes.

Factors Affecting Transmission Performance

1. Source/Target Specifications: During DTS transmission, the load on both the source and target ends will increase. If the load is already high, this can lead to a decline in transmission performance. In scenes where the source or target end has lower specifications, it is recommended to reduce the DTS transmission rate accordingly.

2. DTS Transmission Parameter Configuration: The RPS (Requests Per Second) and concurrency settings for DTS tasks can be flexibly configured. In most cases, the default parameters do not require modification. However, if these settings are not configured properly, they may limit performance.

3. Network Bandwidth and Latency Between Source and Target: The network between the source and target should be stable, with no latency or packet loss, to ensure optimal performance.

When the **Public Network** is selected as an access method, the bandwidth cannot be guaranteed due to the public network environment, which may become a transmission bottleneck. It is recommended to select **Direct Connect** or **VPN Access** to improve network quality.

When the **CCN** is selected as an access method, note that it only provides free bandwidth of up to 10 Kbps between all regions. If you are using DTS to transmit official business data, you will need to pay for higher bandwidth.

Cross-region transmission latency is typically higher than within the same region. The region of the DTS task instance should match the region of the target instance to avoid increased latency. Additionally, for self-hosted databases, it is recommended to select a DTS region that is geographically close to the source for optimal transmission.

4. Source Data Issues: If the data on the source end does not comply with DTS specifications (e.g., long-running SQL queries, tables without primary keys), it may cause the transmission to stall. In such cases, it is recommended to optimize the indexes and queries on the source end. Additionally, large transactions or frequent DDL operations on the source end can result in slower transmission rates during the incremental phase.

DTS Performance Optimization Methods

1. Adjust the DTS task transmission parameters.

Select the specified task, then choose **More > Limit Speed** in the **Operation** column to adjust the parameters.

Adjustable Parameters in the Full Data Phase: The number of concurrent threads and RPS for both the source and target database full data import can be adjusted.

Adjustable Parameters in the Incremental Data Phase: The number of concurrent threads for the target database incremental import can be adjusted.

2. Adjust the DTS task specifications.

Select the specified task, then choose **More > Adjust Specification** in the **Operation** column to adjust the parameters.

When you purchase a DTS task, different specifications correspond to different RPS limits during the incremental phase; for example, in data synchronization tasks, if the write load on the source database is too high and the purchased specification is insufficient, you will need to upgrade the DTS task specifications. For more details, see the [Specification Description](#).

Specifications	Performance Limit (RPS) Reference
micro	1000
small	2000
medium	5000
large	> 5000

3. Pause and then resume the task.

Select the specified task, then choose **More > Pause** in the **Operation** column.

If a temporary spike in source database activity causes a high load on the target end, you can pause the task and resume synchronization after the peak has subsided.

General Principles

When you encounter performance issues, first check the network and configuration factors before adjusting the DTS rate configuration parameters.

DTS transmission performance is influenced by factors such as source/target configurations, DTS task specifications, network bandwidth, and latency. Therefore, simply increasing the number of concurrent threads or RPS may not be effective in improving performance. It is essential to analyze where the performance bottleneck lies and address the specific issue accordingly.

Performance Optimization Steps

1. Monitor the information.

Select the specified task, click **Task ID** to enter the task details page. Then, switch to the relevant tab and click **Monitoring Data** to view the corresponding metric data.

2. If the monitoring data shows high latency, first check for network issues. Latency greater than 3 ms within the same region is considered high, and cross-region latency may be slightly higher.

2.1 When the **Public Network** is selected as an access method, the bandwidth cannot be guaranteed due to the public network environment, which may become a transmission bottleneck. It is recommended to select **Direct Connect** or **VPN Access** to improve network quality.

2.2 When the **CCN** is selected as an access method, note that it only provides free bandwidth of up to 10 Kbps between all regions. For official business data transmission using DTS, you will need to pay for higher bandwidth.

2.3 The DTS task instance should be located in the same region as the target instance; otherwise, latency will increase. For example, if you intend to synchronize data to the Guangzhou region, you should create the synchronization task in the Guangzhou region.

2.4 If you are using a self-hosted database, it is important to select a DTS region that is geographically close to the database for transmission; otherwise, latency will increase. For example, if your self-hosted database is in Beijing and you need to synchronize it to a Tencent Cloud database in the Shanghai region, you should select the region as Beijing -> Shanghai when creating the DTS link. Choosing Shanghai -> Shanghai or Shenzhen -> Shanghai would increase latency.

3. During the **Full Data Phase**, adjust the DTS rate.

Note:

During the **Full Export Phase** of data migration tasks (limited to structure export or export of tables without primary keys) or data synchronization tasks, or when the primary key conflict strategy is set to **Error on Conflict**, adjusting the rate parameters may cause the task to restart from the beginning. Therefore, carefully evaluate before proceeding with the adjustments in sections 3.1 and 3.2.

3.1 Normal export rate: Full export BPS (Bytes Per Second) is higher than 20-50 MB/s, with network latency to the source instance less than 1 ms during full export.

In such cases, the adjustment is generally not necessary. However, if the user needs to increase the rate, it should be done after analyzing the load on the source end.

If the load on the source end is high (e.g., CPU usage greater than 90%), it is not recommended to increase the rate further.

If the load on the source end is low (e.g., CPU usage less than 50%), you can consider increasing the number of concurrent export threads on the source end. In the full data phase, the default RPS is significantly high, so adjusting the RPS is usually unnecessary when trying to improve transmission speed. If increasing the number of export threads does not produce significant results, proceed to step 5.

3.2 Limited export rate: Full export BPS is less than 20 MB/s, and network latency to the source instance during full export is greater than 3 ms.

Analysis should be conducted in conjunction with the load on the source end.

If the load on the source end is high (e.g., CPU usage greater than 90%), it indicates that the source's configuration is relatively low, and the added load from DTS transmission is further limiting performance. It is recommended to reduce the number of concurrent export threads and RPS on the source end. If this does not significantly improve performance, it is recommended to use a replica database for synchronization.

If the load on the source end is low, you can consider increasing the number of concurrent export threads. If this adjustment does not produce significant results, proceed to step 5.

3.3 Normal import rate: Full import BPS is higher than 20-50 MB/s, with network latency to the target instance less than 1 ms during full import.

In such cases, the adjustment is generally not necessary. However, if the user needs to increase the rate, it should be done after considering the resource usage on the target end.

If the load on the target end is high (e.g., CPU usage greater than 90%), it is not recommended to increase the rate further.

If the load on the target end is low (e.g., CPU usage less than 50%), you can consider increasing the number of concurrent import threads on the target end. If this adjustment does not produce significant results, proceed to step 5.

3.4 Limited import rate: Full import BPS is less than 20 MB/s, and network latency to the target instance during full import is greater than 3 ms.

Analysis should be conducted in conjunction with the load on the target end.

If the load on the target end is high (e.g., CPU usage greater than 90%), it indicates that the target's configuration is relatively low, and the added load from DTS transmission is further limiting performance. In this case, it is recommended to reduce the number of concurrent import threads and RPS on the target end.

If the adjustment does not yield significant results, you can use the pause feature to temporarily stop the task transmission. **It is recommended not to pause the task for more than 3 days, as it may not be able to resume transmission afterward.**

If the load on the target end is low, you can consider increasing the number of concurrent import threads. If this adjustment does not produce significant results, proceed to step 5.

4. During the **Incremental Phase**, adjust the DTS rate.

4.1 Check the monitoring data to see if the incremental export RPS from the source instance or the incremental import RPS to the target instance has reached the specified limit. Different task specifications have different RPS limits.

If the incremental RPS has reached the limit, first upgrade the DTS task specifications (done through the **Adjust Specification** operation).

If the incremental RPS is less than 60% of the task specification limit, performance may be restricted. Proceed to step 4.2.

4.2 Check the load on the target end.

If the target end is at its limit (e.g., CPU usage greater than 95%), proceed to step 4.3.

If the target end load is normal (e.g., CPU usage is between 60 and 85%), adjustments are not necessary.

If the target end load is low (e.g., CPU usage is less than 50%), you can consider increasing the DTS rate parameters.

If the target end is a database, you can consider increasing the number of concurrent write threads on the target end.

If this adjustment does not produce significant results, proceed to step 5.

If the target end is Kafka, increasing the number of concurrent write threads may not significantly improve performance. Instead, first increase the bandwidth configuration and the number of partitions for the target CKafka (multiple partitions can be written concurrently, improving the write rate to the target end).

4.3 Check if there are large transactions or frequent DDL operations on the source end.

4.3.1 The source end has large transactions or frequent DDL operations.

This can cause a temporary load on the target end. You can pause the task and resume it after the large transaction or DDL operation has completed. However, it is recommended not to pause the task for more than 3 days, as this may lead to task failure.

4.3.2 The source end has no large transactions or frequent DDL operations.

In this case, the low configuration of the target end, combined with the added load from DTS transmission, is limiting performance. It is recommended to first reduce the number of concurrent import threads on the target end.

If the target end is a database, you can consider reducing the number of concurrent write threads. If this adjustment does not improve performance, you can consider **downgrading** the task configuration.

If the target end is Kafka, you can first increase the bandwidth configuration and the number of partitions on the target CKafka (multiple partitions can be written concurrently, improving the write rate to the target end). Then, reduce the number of concurrent write threads. If these adjustments do not improve performance, you can consider **downgrading** the task configuration.

5. Check the data content on the source end.

5.1 Check if there are long-running SQL queries on the source end.

If there are slow SQL queries running on the source end, it can cause the synchronization rate to slow down. Check for any slow SQL queries on the source end, address them as needed, and consider optimizing indexes to prevent slow SQL from running.

5.2 Check if there is any data on the source end that does not comply with DTS specifications.

If there are tables without primary keys on the source end, large-scale queries on these tables may cause the synchronization process to slow down. It is recommended to add primary keys to these tables or avoid synchronizing tables without primary keys.