

TencentDB for SQL Server

Operation Guide

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

Operation Guide

Constraints and Limits

Usage Specifications and Suggestions

Maintaining Instance

Renaming Instance

Setting Instance Remarks

Setting Instance Tag

Setting Instance Project

Modify Instance-level Character Set Collation

Modify System Time Zone

Setting Instance Maintenance Information

Multi-AZ Disaster Recovery

Restarting Instance

Terminating Instance

Migrating Across AZs

Manual Primary-Secondary Switching

Recycle Bin

Adjusting Instance Configuration

Overview

Adjusting Instance Version

Adjusting Instance Specification

Read-Only Instance

Read-Only Instance Overview

Managing Read-Only Instance

Read-Only Group

Backup node read-only

Network and Security

Switching from Classic Network to VPC

Changing Network

Enabling/Disabling Public Network Address

Enabling Public Network Access Through CLB

CAM

Overview

Authorization Policy Syntax

Authorizable Resource Types

Account Management

- Account Types and Permissions

- Account List

- Creating Account

- Deleting Account

- Modifying Account Permissions

- Creating Admin Account

- Creating Designated Account

- Resetting Password

- Setting Account CAM Verification

Database Management

- Database List

- Creating Database

- Setting Database Permissions

- Cloning Database

- Deleting Database

- Setting Change Data Capture (CDC)

- Setting Change Tracking (CT)

- Shrinking Database

Data Security

- Transparent data encryption

- Access Management

 - CAM Overview

 - Authorization Policy Syntax

 - Authorizable Resource Types

- Configuring Security Groups

- SSL Encryption Settings

Parameter Configuration

- Setting Instance Parameters

- Viewing Parameter Modification Log

Monitoring and Alarms

- Viewing Monitoring Charts

- Monitoring Metrics

- Setting Alarm Policies

- Setting Alarm Notification

- Viewing Alarm Records

Backup and Restoration

- Backup

Backup Overview

Backup Fees

Setting Automatic Backup

Setting Non-Archive Backup Retention

Setting Archive Backup Retention

Creating Manual Backup

Setting Backup Task

Cross-Region Backup

Viewing Backup List

Downloading Backup

Deleting Manual Backup

Viewing Backup Space

Rolling back

Rollback Scheme Overview

Rolling back to the current instance

Rolling back to an existing instance (same region)

Rollback to an Existing Instance (Cross-Region)

Viewing Rollback Task List

Log Management

Querying and Downloading Blocking and Deadlock Events

Querying and Downloading Slow Query Log

Querying primary and standby switch logs

Manual Creation of the Latest Blocking and Deadlock Events

Publish-Subscribe

Overview of Publish-Subscribe

Managing Publish-Subscribe

SSIS

Overview

Purchasing Business Intelligence Server

Creating Windows Authentication Account

Adding Flat File

Interconnecting Source/Target Instances and Business Intelligence Server

Deploying Project

Data Migration (New)

Data Migration Solution Overview

Cold Backup Migration

Using DTS to Migrate Data

Instructions

Migration Operation Guide

Cloud Database Across-Account Migration Guide

Pre-Validation Failure Handling

Migration Parameter Check

Migration Network Check

Source Database Connectivity Check

Target Database Connectivity Check

Disk Space Check

Character Set check

Database Version Check

User Permission Check

Source Database Existence Check

Target Database Existence Check

Error Handling

Data Migration (Legacy)

Migrating Data from CVM-based SQL Server Instance

Migrating Data from SQL Server Instance Purchased at Other Cloud Service Provider or Built In-house

Same-Version Migration

Cross-version Migration

Operation Guide

Constraints and Limits

Last updated : 2024-08-02 17:36:10

TencentDB for SQL Server service only offers instances with bundled licenses, meaning that upon creation, each instance is endowed with the corresponding version's Microsoft SQL Server software licensing authorization. It does not support the provision of licenses brought by users themselves.

In order to ensure the stability and security of instances, TencentDB for SQL Server has certain restrictions on usage as detailed below.

TencentDB for SQL Server is available in two editions: single-node edition (formerly Basic Edition) and two-node edition (formerly High Availability/Cluster Edition), each with its own set of features. For more information, see [Features and Differences](#).

Note:

TencentDB for SQL Server does not support access to external applications through the database instance; it only permits access to the database instance via external applications.

If you have other questions about usage restrictions, [submit a ticket](#) for assistance.

Feature	Two-Node Edition	Single-Node Edition
Database version	2008 R2 Enterprise 2012 Enterprise 2014 Enterprise 2016 Enterprise 2017 Enterprise 2019 Enterprise 2022 Enterprise	
Maximum number of databases (subject to the number of instance CPU cores as described in Constraints and Limits > Database quantity)	2008 R2 version: 300 2012, 2014, and 2016 versions: 300 2017, 2019, and 2022 versions: 100	400
Maximum number of database accounts	Unlimited	Unlimited
Database creation	Supported	Supported
User/Login account creation and deletion	Supported	Supported
SA account creation	Not supported	Supported

Database authorization	Supported	Supported
Database-level DDL trigger	Supported	Supported
Thread killing permission	Supported	Supported
SQL profiler	Supported	Supported
Publish/Subscribe	Supported	Not supported
Optimization advisor	Not supported	Supported
Linked server	For a detailed explanation, please refer to Limitations of the Linked Server Features .	
Distributed transaction	Only supported by two-node local disk instances	Not supported
Change data capture (CDC)	Supported	Supported
Change tracking (CT)	Supported	Supported
Windows domain account login	Not supported	Not supported
Email	Not supported	Not supported
SQL Server Integration Services (SSIS)	Supported	Supported
SQL Server Analysis Services (SSAS)	Not supported	Not supported
SQL Server Reporting Services (SSRS)	Not supported	Not supported
R language service	Not supported	Not supported
Common Language Runtime (CLR) integration	Not supported	Not supported
Async messaging	Not supported	Not supported
Policy management	Not supported	Not supported

Database quantity

Note:

For a two-node (formerly High Availability/Cluster Edition) instance, if you set `max worker threads` to the default value 0, you can create no more than 100 databases. To create more databases, you must set this parameter to 20,000 as instructed in [Setting Instance Parameters](#).

If the instance only has one CPU core, we recommend that you keep the database quantity limit at 70 to guarantee instance stability.

SQL Server 2008 R2 Enterprise instances don't support lifting the database quantity limit, which is 70. The limit in other SQL Server instances is subject to the number of instance CPU cores as calculated below:

Two-node (formerly High Availability Edition)

2012 Standard/Enterprise

2014 Standard/Enterprise

2016 Standard/Enterprise

Maximum number of databases:

$$\min\{80 + \sqrt{\text{CPU core quantity}} * 40, 300\} \text{ (Extract the square root and round it to one decimal place)}$$

Extract the square root of the CPU core quantity, round it to one decimal place, multiply the result by 40, and add the product to 80 to get the value X. The smaller value between X and 300 is the maximum number of databases. For example, you can create up to 160 databases in a 4-core 16 GB MEM SQL Server 2014 Enterprise instance.

Two-node (formerly Cluster Edition)

2017 Enterprise

2019 Enterprise

2022 Enterprise

The maximum number of databases is related to the instance CPU. The maximum number of databases for instances with 8-core CPU or less is 80, and the maximum number of databases for instances with 8-core CPU or above is 100.

Single-node (formerly Basic Edition)

2008 R2 Enterprise

2012 Enterprise

2014 Enterprise

2016 Enterprise

2017 Enterprise

2019 Enterprise

2022 Enterprise

Maximum number of databases:

$$\min\left\{\left\lfloor\sqrt{\text{CPU core quantity}}\right\rfloor * 100, 400\right\}$$

Extract the square root of the CPU core quantity, round it down to the nearest integer, and multiply the result by 100 to get the value N. The smaller value between N and 400 is the maximum number of databases. For example, you can create up to 200 databases in a 4-core 16 GB MEM SQL Server 2017 Enterprise instance.

Table of instance CPU core quantity and corresponding database quantity limit

Database quantity limit for two-node edition (formerly High Availability/Cluster Edition)

Database quantity limit for single-node edition (formerly Basic Edition)

CPU cores	Two-node 2008 R2/2012/2014/2016 Enterprise	Two-node 2017/2019/2022 Enterprise
1	70	70
2	136	176
4	160	200
8	193	233
12	218	258
16	240	280
24	275	315
32	300	340
48	300	340
64	300	340
96	300	340

Number of CPU Cores	Single Node 2008 R2 Enterprise Edition	Single-node 2012/2014/2016/2017/2019/2022 Enterprise Edition
2	80	100
4	80	200
8	80	200
12	80	300

16	80	400
24	80	400
32	80	400
48	80	400
64	80	400

Linked Server Limitations

Linked servers are not supported between other clouds or self-built environments and Tencent Cloud and are only supported for a Tencent Cloud private network. However, within the TencentDB for SQL Server instance architecture, there are the following restrictions:

1. Cross-Architecture Scenarios

Cross-Architecture Direction	Description
Two-node instance → single-node instance	Not supported.
Single-node instance → two-node instance	Not supported.
Local disk instance → cloud disk instance	Not supported.
Cloud disk instance → local disk instance	Not supported.

2. Between Two-Node Instances

Cross-Architecture Direction	Description
Two-node local disk instance → two-node local disk instance	Interconnection is supported between domestic regions, and interconnection is supported between overseas regions, but not between domestic and overseas regions.
Two-node cloud disk instance → two-node cloud disk instance	Only interconnection within the same region is supported.

3. Between Single Node Instances

Interconnection is supported only when the source and target are the same instance.

Usage Specifications and Suggestions

Last updated : 2024-01-18 17:20:33

This document describes the usage specifications and suggestions for TencentDB for SQL Server.

Purpose

To standardize the management and maintenance of TencentDB for SQL Server to avoid unavailability and other issues caused by improper operations.

To provide guidance for database developers on how to write SQL statements to ensure optimal performance of TencentDB for SQL Server.

Suggestions for instance specification

Do not use instances with the 1-core specification in the production environment, as this specification is only suitable for feature testing.

Use instances with the 2-core or higher specification in the production environment. As SQL Server runs on Windows, the engine and the system require many resources. Therefore, the 1-core specification isn't suitable for sustaining the production business, and problems such as low system memory and system lags may occur after long-time operations on this specification.

Suggestions for instance selection

Use two-node (formerly High Availability/Cluster Edition) primary/replica instances. Compared with a single-node (formerly Basic Edition) instance, two-node instances can greatly improve the availability and reliability of the production business.

If your business has few write requests but massive read requests and you need to add read-only instances, use SQL Server 2017/2019 two-node instances to sync data more efficiently and stably.

Select multi-AZ deployment for Dual-Server High Availability/Cluster Edition instances to implement AZ-level disaster recovery.

Suggestions for database connection

Use `ip,port` to connect to a TencentDB for SQL Server instance. Note that the IP and port are separated by a comma.

Do not use the server name for connection. Your application should better have a reconnection mechanism for database connection, so it can reconnect to databases in time through retries in case of database failure or disconnection.

Permission management specifications

To ensure the stability and security, permission restrictions are imposed on `sysadmin` and `shutdown` in TencentDB for SQL Server. The following errors may occur when you run certain statements:

```
User does not have permission to perform this action.
```

```
You do not have permission to run the RECONFIGURE statement.
```

Solution: Modify parameters, manage databases and users, and restore backups in the console.

Grant permissions on demand. You only need to grant general applications the read/write permissions for specified databases.

Grant permissions to users of general applications at the database level.

Allow authorized users to access TencentDB for SQL Server only from specific IPs or IP ranges. This can be achieved by configuring security groups in the console as prompted.

Separate management, development, and application accounts. Avoid using admin accounts for development or business operations.

Specifications and suggestions for routine operations

Too many databases will compromise the TencentDB for SQL Server instance performance and occupy more resources such as worker threads. If the limit on the number of created instances is exceeded, primary/replica sync exceptions may occur. We recommend that you keep the number of databases created in a single instance below the upper limit, which is subject to the CPU core quantity of the instance. For more information, see [Constraints and Limits > Database quantity](#).

The database name can contain up to 64 digits, letters, and underscores.

For enhanced instance security, do not use weak passwords. Perform account and database management operations in the console in general cases.

For login over the private network, make sure that the CVM instance of the client and the TencentDB for SQL Server instance are in the same VPC in the same region under the same account.

Applications should not rely on `sysadmin` permissions to use databases. Accounts with the `sysadmin` role have the super admin permissions. If they are used improperly, the database security and stability will be compromised. Therefore, TencentDB doesn't grant the super admin permissions by default.

Check the database size and shrink databases promptly. If a database is used for a long time, the used physical space may not be released in time, and you need to shrink the database to release such space. Check the log file size and physical file size frequently, so that you can shrink databases during off-peak hours when you find that the file size increases rapidly.

After long-term operations, instances may suffer from a performance drop. Therefore, restart instances at least once every three months during off-peak hours.

Do not create tables or write data in system databases. Store data in created custom databases. Although permissions to use system databases are granted, any data stored in system databases is insecure.

Do not set databases to the single-user mode. In this mode, only one session is allowed to access databases, causing TencentDB Ops problems. If the single-user mode is set, roll back to the multi-user mode promptly.

Extended events are used for slow log collection. This is a lightweight tracking method that basically has no impact on instances.

Reindex database regularly. After a database runs for a long time, it may generate many index fragments, which compromise the database access performance. Therefore, you need to regularly reindex databases by creating SQL agent jobs, preferably once every month.

Update the statistics regularly by creating SQL agent jobs, preferably once every week. This helps guarantee the performance.

Set the maximum concurrency, which determines the business CPU utilization.

Perform backup and restoration operations through the console or APIs rather than SSMS or SQL statements. For more information on backup and restoration methods and database migration to the cloud, see [Cold Backup Migration](#).

Do not set the database recovery mode to "simple". Use "full" instead.

If the "simple" recovery mode is used, incremental backup won't be implemented on the database, so the database cannot be restored to the specified time point.

For two-node (formerly High Availability/Cluster Edition) instances, after you set the database recovery mode to "simple", the database cannot establish replication relationships and thus cannot support primary/replica switch or specification change.

Therefore, use the "simple" recovery mode with caution.

Avoid performing DDL operations during peak hours.

Avoid performing batch operations during peak hours. To delete an entire table, use `TRUNCATE` or `DROP` during off-peak hours.

Avoid running an instance for multiple businesses to minimize the risk of mutual interference between businesses due to high coupling.

Avoid using automatic transaction committing and develop a habit of using `begin tran;` for online operations, which helps minimize the risk of data loss caused by misoperations. In case of a misoperation, you can use the rollback feature of TencentDB for SQL Server for data restoration. After a transaction begins, commit it in time to avoid instance blocking.

Perform database operations in the console rather than on the SSMS client.

Estimate the resources required in advance and optimize the instances for promotional campaigns of your business. In case of a great demand for resources, contact your Tencent Cloud sales rep timely.

Suggestions for using DTS for database migration

Check the following before migrating a database to the cloud:

Version numbers of source and target databases. The target database must be on a version later than or equal to the source database. For example, if the source database is on v2016, the target database can only be on v2016, v2017, or v2019.

Architecture versions of source and target databases. If the source instance is a self-built database in a local IDC, CVM instance, or cloud server in another cloud vendor, or is a cloud SQL Server instance in another cloud vendor, you can migrate it to a TencentDB for SQL Server single-node (formerly Basic Edition) instance or two-node (formerly High Availability/Cluster Edition) instance on any architecture version. If the source instance is a TencentDB for SQL Server two-node instance, it cannot be migrated to a single-node instance through DTS. If the source instance is a TencentDB for SQL Server single-node instance, it can be migrated to a two-node instance through DTS.

Network connectivity between source and target databases. The source and target databases must be connected. The server where the source database resides must have enough outbound bandwidth; otherwise, the migration efficiency will be affected.

Names of source and target databases. The source and target instances cannot contain databases with the same name.

Account permissions of the source database. You need to change to `local` for SQL service startup in the source database. The source database account is unrestricted but needs to have the `sysadmin` permissions.

Account permissions of the target database. The target database needs to have an account with admin permissions for migration.

Ports of the source database. The source database needs to open port 1433, and the service where the source database is located must open the file sharing port 445 for Windows server sharing.

Recovery mode of the source database. The source database must be set to "full recovery mode", and we recommend that you make a full backup before migration.

Local disk space of the source database. The local disk space of the source database must be large enough, so that the remaining free space can fit the size of the database to be migrated.

Disk space of the target database. The disk space of the target database must be at least 1.5 times the size of the source database.

Status of the target database. The target database cannot have access requests or active businesses; otherwise, the migration will fail.

You need to keep the following operation limits in mind when migrating data to the cloud:

Only one migration task can be initiated at any time for the same source instance.

Only database-level migration is supported, that is, all objects in the database must be migrated together. Single-table migration is not supported.

Logins, jobs, triggers, and database links (link servers) at the instance level cannot be migrated.

Do not modify or delete user information (including username, password, and permissions) in the source and target databases and port numbers during migration; otherwise, the migration task will fail.

Do not perform transaction log backup during incremental sync; otherwise, the transaction log will be truncated and become discontinuous.

If you only perform full data migration, do not write new data into the source database 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 that you select full + incremental data migration.

For full + incremental data migration, after you click **Complete** and the task status becomes **Completed**, do not write new data to the source database. We recommend that you stop writing for two minutes; otherwise, the data in the source and target databases may be inconsistent.

Check the following after migrating data to the cloud:

Permission completeness. Permissions will affect operations performed on the database. The migration only restores data. To restore other service-level permissions, such as database users and login usernames, you need to create them again and associate them with database accounts.

Reindexing. As the physical environment of the data files changes, database indexes will become invalid, and you need to create indexes again; otherwise, the database performance may be significantly compromised.

Instance-level objects such as logins, jobs, triggers, and database links (link servers). You need to create them again after the migration is completed.

Database and table design specifications

Note

TencentDB for SQL Server versions earlier than 2014 don't support memory-optimized tables. If you need to use this type of tables, we recommend that you use TencentDB for Redis and Memcached.

Follow the third normal form (3NF) when creating tables and specify the primary key for each table. Even if you can't select an appropriate column as the primary key, you still need to select one.

Define fields as NOT NULL and set default values. NULL fields will cause unavailability of indexes, thus bringing problems to SQL development. NULL calculation can only be implemented based on IS NULL and IS NOT NULL.

Suggestions:

Plan the resources used by databases reasonably based on business scenario analysis and estimation of data access (including database read/write QPS, TPS, and storage). You can also configure various monitoring metrics for TencentDB for SQL Server in the Tencent Cloud Observability Platform (TCOP) console.

Put the tables for the same type of businesses into one database when building databases. Do not perform cross-database correlation operations in programs, as doing so will affect subsequent quick rollbacks.

Use the same character set to avoid garbled text caused by conflicts between different character sets.

Use the `DECIMAL` type to store decimal values. The `FLOAT` and `DOUBLE` types have insufficient precision, especially for businesses involving money.

Do not use the `TEXT` or `BLOB` type to store a large quantity of text, binary data, images, files, and other contents in a database; instead, save such data as local disk files and only store their index information in the database.

Avoid using foreign keys. We recommend that you implement the foreign key logic at the application layer. Foreign key and cascade update are not suitable for high-concurrency scenarios, because they may reduce the insertion performance and lead to deadlock in case of high concurrency.

Reduce the coupling of business logic and data storage, use databases mainly for storing data, implement business logic at the application layer, and minimize the use of instance-level triggers, link servers, jobs, and other advanced features, given their poor portability and scalability. If such objects exist in an instance, you need to migrate them to the new instance manually after data migration.

If you don't have a significant increase in business volume in the near future, do not use partitioned tables, which are mainly used for archive management. Partitioned tables have no obvious improvement on the performance if most queries in your business don't involve partition fields.

Purchase read-only instances to implement read/write separation at the database level for business scenarios with a high read load and low requirement for consistency (where a data delay within seconds is acceptable).

Index design specifications

Note

Do not create indexes on the columns that are updated frequently and have a lower selectivity. Recording updates will change the B+ tree, so creating indexes on frequently updated fields may greatly compromise the database performance.

Put the column with the highest selectivity on the far left when creating a composite index; for example, in `select xxx where a = x and b = x;`, if a and b are used together to create a composite index and a has a higher selectivity, then the composite index should be created as `idx_ab(a,b)`. If `Not Equal To` and `Equal To` conditions are used at the same time, the column with the `Equal To` condition must be put first; for example, in `where a xxx and b = xxx`, b must be placed on the far left even if a has a higher selectivity, because a will not be used in the query.

Suggestions:

Use no more than five indexes in a single table and no more than five fields in a single index. Too many indexes may affect the filtering, occupy much more capacity, and consume more resources for management.

Create indexes on the columns that are used for SQL filtering most frequently with a small number of duplicate values. It is meaningless to create indexes on a column not involved in SQL filtering. The higher the uniqueness of a field, the better the index filtering effect.

Avoid using redundant indexes. If both index (a,b) and index (a) exist, (a) is considered a redundant index. If the query filtering is based on column a, the index (a,b) is sufficient.

Use `INCLUDE index` reasonably to reduce I/O overheads. Commonly used columns should be placed on the left, and columns that will not be used as query conditions can be placed in `INCLUDE`.

SQL statement writing specifications

Note

For `UPDATE` and `DELETE`, use `WHERE` for exact match. To delete a large amount of data, delete the data in batches during off-peak hours.

When using `INSERT INTO t_xxx VALUES (xxx)`, explicitly specify the column attributes to be inserted to prevent data errors caused by changes in the table structure.

The following are common causes of invalid indexes in SQL statements:

Implicit type conversion; for example, if the type of index a is `VARCHAR` and the SQL statement is `where a = 1`, then `VARCHAR` is changed to `INT`.

Math calculations and functions are performed on the index columns; for example, date column is formatted using a function.

Multiple columns with different sorting orders; for example, the index is (a,b), but the SQL statement is `order by a b desc`.

Use the `WHERE` condition for exact match to avoid fuzzy match of conditions and batch matches of `IN` and `NOT IN`.

Suggestions:

Ensure query on demand and reject `select *` to avoid the following problems:

The covering index does not work and the problem of `TABLE ACCESS BY INDEX ROWID` occurs, which leads to extra I/O overheads.

There is an extra memory load. A large amount of cold data is imported to the cache, which may reduce the query hit rate.

There are extra overheads in network transfer.

Instance blocking or primary/replica delay may occur. To prevent this, avoid using large transactions and split a large transaction into multiple small ones.

Unnecessary lock waits may occur. To prevent this, commit transactions in the business code timely.

Minimize the use of `JOIN` operations on multiple tables and big tables. When joining two tables, use the smaller one as the driving table and select indexed columns with the same character set for join.

Description

It is difficult to completely avoid the aforementioned issues. The solution is to set the aforementioned conditions as secondary filtering conditions for indexes rather than as primary filtering conditions.

If a large number of full-table scans is found in the monitoring data, you can download slow log files in the console for analysis.

Perform the required SQL audit before a business is released. In routine Ops work, download slow query logs regularly for targeted optimization.

Maintaining Instance

Renaming Instance

Last updated : 2024-01-18 17:20:33


TencentDB for SQL Server instances are differentiated by name, which can be renamed in the console. This document describes how to rename an instance in the console.

Instance ID/name description

Both primary and read-only instances can be renamed.

An instance name can contain up to 60 Chinese characters, letters, digits, or underscores.

An instance ID supports quick copy.




Renaming an instance in the instance list

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click



after **Instance Name**.



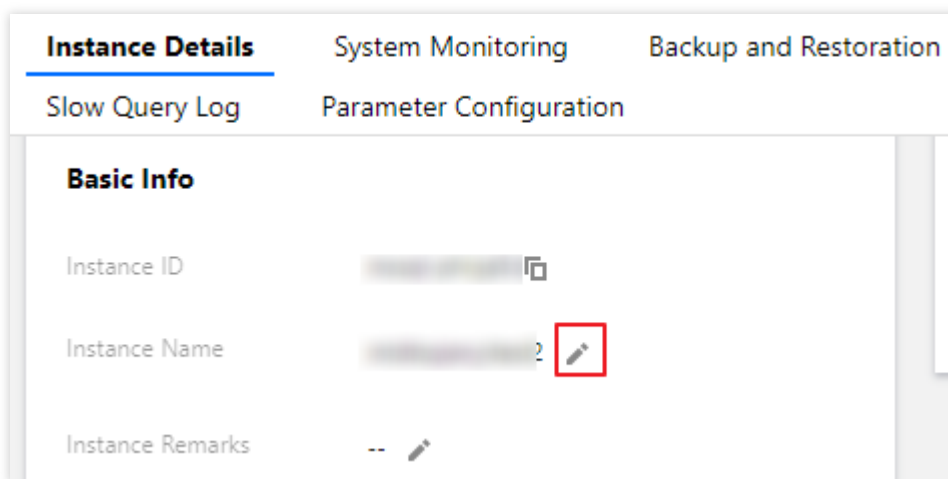
3. In the pop-up window, enter a new name and click **OK**.

Renaming an instance on the instance details page

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, select **Instance Details** > **Basic Info** and click



after **Instance Name**.



4. In the pop-up window, enter a new name and click **OK**.

Setting Instance Remarks

Last updated : 2024-01-18 17:20:33

Overview

As your business grows, you will need to manage an increasing number of devices and instances, thus it is in your great interest to properly categorize your resources. To make it simple for you to manage and identify your instances, TencentDB for SQL Server supports adding remarks to instances in the console.

Notes


Both primary and read-only instances support remarks.

Instance remarks can contain up to 200 Chinese characters, letters, digits, or underscores.

You can edit, modify, and delete instance remarks.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. In **Instance Details > Basic Info** on the instance management page, click



after **Instance Remarks**.

4. In the pop-up window, edit the remarks and click **OK**.

Modify Instance Remarks✕

Instance Remarks

Up to 200 characters

0/200 (It can contain up to 200 letters, digits, or underscores.)

OK

Cancel

Setting Instance Tag

Last updated : 2024-01-18 17:20:33

Tag Overview

Tags are key-value pairs provided by Tencent Cloud to easily identify resources. For more information, see [Tag Overview](#).

You can use tags to categorize and manage TencentDB for SQL Server resources by various metrics such as business, purpose, and owner. You can also easily find a resource by its tag. In Tencent Cloud, the tag key-value pairs have no semantic meaning and are strictly parsed and matched as strings. To use tags, pay attention to the [use limits](#) first.

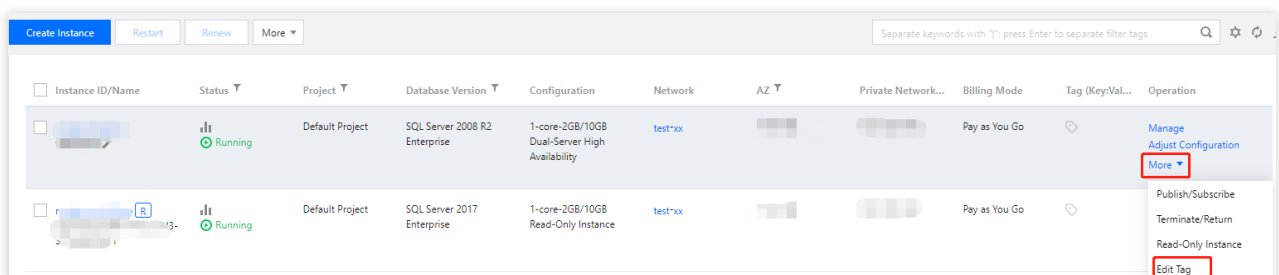
Creating Tags

1. Log in to the [Tag console](#).
2. Click **Tag List** on the left sidebar to enter the tag list page.
3. Click **Create Tag**. In the pop-up window, click **Add tag key** to create a tag or select an existing tag and add a tag value to it. You can create multiple tags at a time.
4. Click **OK**.

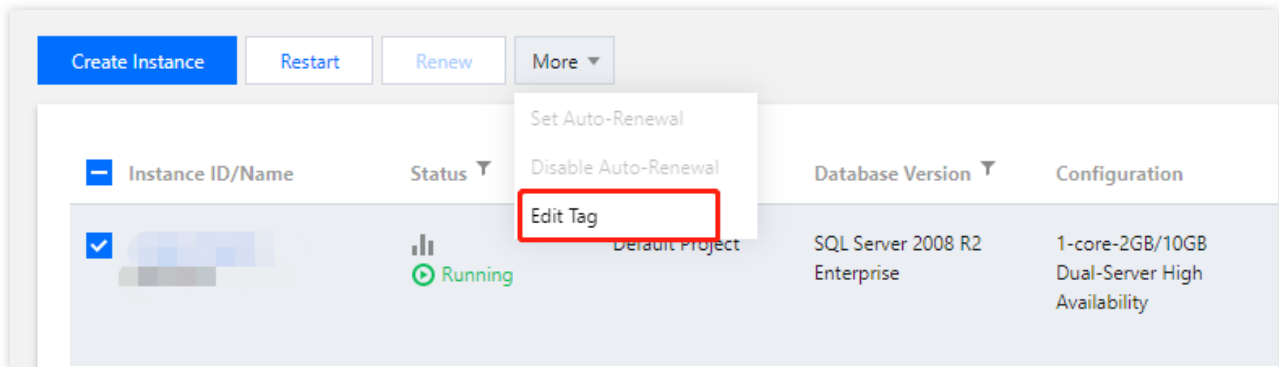
Tagging a TencentDB for SQL Server Instance

Tagging in the instance list

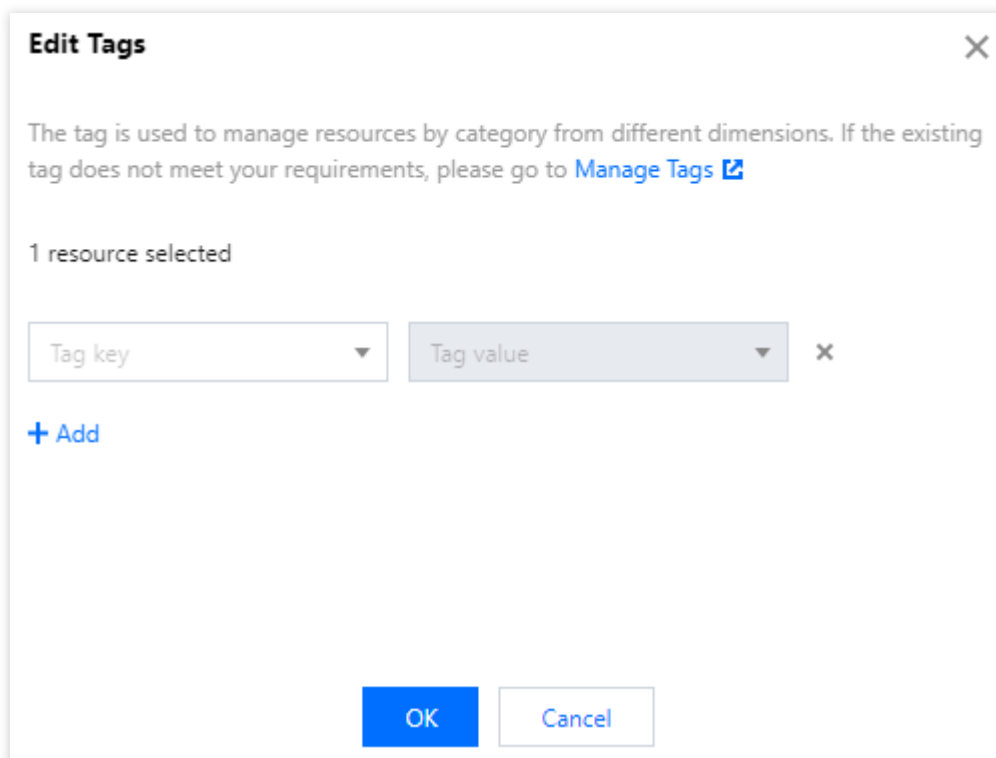
1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and select **More > Edit Tag** in its **Operation** column.



To batch edit tags, select target instances and click **More > Edit Tag** at the top.



3. In the pop-up window, set **Tag Key** and **Tag Value** and click **OK**.



Tagging on the instance details page

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, select the **Instance Details** tab and click



after **Tag** in **Basic Info**.

4. In the pop-up window, set **Tag Key** and **Tag Value** and click **OK**.

Setting Instance Project

Last updated : 2024-01-18 17:20:33

TencentDB for SQL Server supports assigning instances to different projects for management.

Log in to the [TencentDB for SQL Server Console](#). In the instance list, click an instance ID to enter the details page, and click **Switch to another project** in "Project" to assign the instance to a project.

Read-only instances are associated instances of the master instance and should be in the same project as the master instance.

Assigning and moving database instances across projects will not affect the services provided by the instances.

You need to specify a project to which a new instance belongs when purchasing it. The default project is **Default Project**.

Assigned instances can be reassigned to other projects through the **Switch to another project** feature.

Modify Instance-level Character Set Collation

Last updated : 2024-01-18 17:20:33

TencentDB for SQL Server allows users to modify the instance-level character set collation when creating an instance and meets demand of some users to set the default character set themselves. This operation can be performed only on the purchase page. The character set of an instance provides a collation for system data, namely the case sensitivity and accent sensitivity. Selecting a collation for database will affect the results of relevant database operations.

This document describes how to modify the character set collation.

Note:

Dual-Node Local Disk Instance:

As **modifying the instance-level character set collation requires to separately configure physical machine resources**, [submit a ticket](#) for assistance and specify the desired character set collation if any modifications are needed before purchase.

For instances with modified character set collation, if **a subsequent expansion involves data migration**, [submit a ticket](#) for assistance.

The default collation of instance character set is Chinese_PRC_CI_AS.

Single-Node Cloud Disk/ Dual-Node Cloud Disk Instances:

The instance-level character set collation can be directly modified on the purchase page.

Prerequisites

If your instance is a dual-node local disk instance, [submit a ticket](#) before modifying the character set collation.

If you purchase a single-node cloud disk or dual-node cloud disk, you can directly modify the character set collation on purchase page with no need for application.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Create Database Instance**.
3. On the instance purchase page, set the **Character Set Collation** under the **System Time Zone** parameter.

System Time	China Standard Time	▼
Zone	Only 'China Standard Time' (UTC+8) is supported by default. To modify it, submit a ticket for assistance.	
Character Set	Chinese_PRC_CI_AS	▼
Collation ?	Only 'Chinese_PRC_CI_AS' is supported by default. To modify it, submit a ticket for assistance.	

- For more information on how to configure other parameters on the purchase page, see [Creating TencentDB for SQL Server Instance](#). After setting all the parameters, click **Buy Now**.
- You can query the configured character set collation under **Basic Info** on the **Instance Details** tab.

Basic Info

Instance ID	...	🔗
Instance Name	...	✎
Instance Remarks	--	✎
Status	Running	
Version	SQL Server 2012 Enterprise	
Architecture	Dual-Server High Availability	
Memory	2GB	
Total Capacity	10GB	
Region	...	
	Migrate across AZs	
AZ	...	
Network	...	
	Change Network	
System Time Zone	China Standard Time	
System Character Set Collation	Chinese_PRC_CI_AS	-

Instance Architecture Diagram

Read-only Instance Info

Collation description

Collation Option	Description
Case Sensitive (_CS)	Case sensitive. This option indicates that lowercase letters will precede their corresponding uppercase letters during sorting.
Case Insensitive (_CI)	Case insensitive. This option indicates that the sorting will be case-insensitive.
Accent Sensitive (_AS)	Accent sensitive. This option indicates that the sorting will be accent-sensitive; for example, "a" and "á" are different characters.
Accent Insensitive (_AI)	Accent insensitive.
Binary Value of Characters (_BIN)	This option indicates that the sorting is based on the binary values in character set, without regard to differences in phonetics, glyph, case, and accent marks.

Instance character set suffix description

Instance Character Set Suffix	Description
_CI_AI	Case insensitive and accent insensitive.
_CI_AS	Case insensitive and accent sensitive.
_CS_AI	Case sensitive and accent insensitive.
_CS_AS	Case sensitive and accent sensitive.

Collation by server level

The following table lists the default collations defined by certain OS region settings.

ID	Character Set Collation
0	Chinese_PRC_CI_AS
1	SQL_Latin1_General_CP1_CI_AI
2	Chinese_PRC_CS_AS
3	Latin1_General_CI_AS
4	Latin1_General_100_CI_AS

5	SQL_Latin1_General_CP1_CI_AS
6	SQL_Latin1_General_CP1250_CI_AS
7	SQL_Latin1_General_CP1251_CI_AS
8	SQL_Latin1_General_CP1253_CI_AI
9	SQL_Latin1_General_CP1253_CI_AS
10	SQL_Latin1_General_CP1254_CI_AS
11	SQL_Latin1_General_CP1255_CI_AS
12	SQL_Latin1_General_CP1256_CI_AS
13	SQL_Latin1_General_CP1257_CI_AS
14	SQL_Latin1_General_CP437_BIN
15	SQL_Latin1_General_CP437_BIN2
16	SQL_Latin1_General_CP437_CI_AI
17	SQL_Latin1_General_CP437_CI_AS
18	SQL_Latin1_General_CP850_BIN
19	SQL_Latin1_General_CP850_BIN2
20	SQL_Latin1_General_CP850_CI_AI
21	SQL_Latin1_General_CP850_CI_AS
22	SQL_Latin1_General_Pref_CP1_CI_AS
23	SQL_Latin1_General_Pref_CP437_CI_AS
24	SQL_Latin1_General_Pref_CP850_CI_AS
25	Albanian_CI_AI
26	Albanian_100_CI_AI
27	Arabic_CI_AI
28	Arabic_100_CI_AI
29	Assamese_100_CI_AI

30	Azeri_Cyrillic_100_CI_AI
31	Azeri_Latin_100_CI_AI
32	Bashkir_100_CI_AI
33	Bengali_100_CI_AI
34	Bosnian_Cyrillic_100_CI_AI
35	Bosnian_Latin_100_CI_AI
36	Breton_100_CI_AI
37	Chinese_Hong_Kong_Stroke_90_CI_AI
38	Chinese_PRC_CI_AI
39	Chinese_PRC_90_CI_AI
40	Chinese_PRC_Stroke_CI_AI
41	Chinese_PRC_Stroke_90_CI_AI
42	Chinese_Simplified_Pinyin_100_CI_AI
43	Chinese_Simplified_Stroke_Order_100_CI_AI
44	Chinese_Taiwan_Bopomofo_CI_AI
45	Chinese_Taiwan_Bopomofo_90_CI_AI
46	Chinese_Taiwan_Stroke_CI_AI
47	Chinese_Taiwan_Stroke_90_CI_AI
48	Chinese_Traditional_Bopomofo_100_CI_AI
49	Chinese_Traditional_Pinyin_100_CI_AI
50	Chinese_Traditional_Stroke_Count_100_CI_AI
51	Chinese_Traditional_Stroke_Order_100_CI_AI
52	Corsican_100_CI_AI
53	Croatian_CI_AI
54	Croatian_100_CI_AI

55	Cyrillic_General_CI_AI
56	Cyrillic_General_100_CI_AI
57	Czech_CI_AI
58	Czech_100_CI_AI
59	Danish_Greenlandic_100_CI_AI
60	Danish_Norwegian_CI_AI
61	Dari_100_CI_AI
62	Divehi_90_CI_AI
63	Divehi_100_CI_AI
64	Estonian_CI_AI
65	Estonian_100_CI_AI
66	Finnish_Swedish_CI_AI
67	Finnish_Swedish_100_CI_AI
68	French_CI_AI
69	French_100_CI_AI
70	Frisian_100_CI_AI
71	Georgian_Modern_Sort_CI_AI
72	Georgian_Modern_Sort_100_CI_AI
73	German_PhoneBook_CI_AI
74	German_PhoneBook_100_CI_AI
75	Greek_CI_AI
76	Greek_100_CI_AI
77	Hebrew_CI_AI
78	Hebrew_100_CI_AI
79	Hungarian_CI_AI

80	Hungarian_100_CI_AI
81	Hungarian_Technical_CI_AI
82	Hungarian_Technical_100_CI_AI
83	Icelandic_CI_AI
84	Icelandic_100_CI_AI
85	Indic_General_90_CI_AI
86	Indic_General_100_CI_AI
87	Japanese_CI_AI
88	Japanese_90_CI_AI
89	Japanese_Bushu_Kakusu_100_CI_AI
90	Japanese_Bushu_Kakusu_140_CI_AI
91	Japanese_Unicode_CI_AI
92	Japanese_XJIS_100_CI_AI
93	Japanese_XJIS_140_CI_AI
94	Kazakh_90_CI_AI
95	Kazakh_100_CI_AI
96	Khmer_100_CI_AI
97	Korean_90_CI_AI
98	Korean_100_CI_AI
99	Korean_Wansung_CI_AI
100	Lao_100_CI_AI
101	Latin1_General_CI_AI
102	Latin1_General_100_CI_AI
103	Latvian_CI_AI
104	Latvian_100_CI_AI

105	Lithuanian_CI_AI
106	Lithuanian_100_CI_AI
107	Macedonian_FYROM_90_CI_AI
108	Macedonian_FYROM_100_CI_AI
109	Maltese_100_CI_AI
110	Maori_100_CI_AI
111	Mapudungan_100_CI_AI
112	Modern_Spanish_CI_AI
113	Modern_Spanish_100_CI_AI
114	Mohawk_100_CI_AI
115	Nepali_100_CI_AI
116	Norwegian_100_CI_AI
117	Pashto_100_CI_AI
118	Persian_100_CI_AI
119	Polish_CI_AI
120	Polish_100_CI_AI
121	Romanian_CI_AI
122	Romanian_100_CI_AI
123	Romansh_100_CI_AI
124	Sami_Norway_100_CI_AI
125	Sami_Sweden_Finland_100_CI_AI
126	Serbian_Cyrillic_100_CI_AI
127	Serbian_Latin_100_CI_AI
128	Slovak_CI_AI
129	Slovak_100_CI_AI

130	Slovenian_CI_AI
131	Slovenian_100_CI_AI
132	Syriac_90_CI_AI
133	Syriac_100_CI_AI
134	Tamazight_100_CI_AI
135	Tatar_90_CI_AI
136	Tatar_100_CI_AI
137	Thai_CI_AI
138	Thai_100_CI_AI
139	Tibetan_100_CI_AI
140	Traditional_Spanish_CI_AI
141	Traditional_Spanish_100_CI_AI
142	Turkish_CI_AI
143	Turkish_100_CI_AI
144	Turkmen_100_CI_AI
145	Uighur_100_CI_AI
146	Ukrainian_CI_AI
147	Ukrainian_100_CI_AI
148	Upper_Sorbian_100_CI_AI
149	Urdu_100_CI_AI
150	Uzbek_Latin_90_CI_AI
151	Uzbek_Latin_100_CI_AI
152	Vietnamese_CI_AI
153	Vietnamese_100_CI_AI
154	Welsh_100_CI_AI

155	Yakut_100_CI_AI
156	SQL_AltDiction_CP850_CI_AI
157	Korean_Wansung_CI_AS
158	Chinese_PRC_BIN

Modify System Time Zone

Last updated : 2024-01-18 17:20:33

This document describes how to modify the system time zone when creating a SQL Server instance.

Note :

Dual-Node Local Disk Instance:

As **modifying the system time zone requires to separately configure physical machine resources**, [submit a ticket](#) for assistance and specify the desired system time zone if any modifications are needed before purchase.

Modifications are only supported when purchasing instances of the 90-core 720GB specification. If you need to modify the system time zone, it is recommended to directly purchase the cloud disk version of the dual-node instance, which can be directly modified on the purchase page without specification restrictions.

For instances with modified system time zone, the actual **data stored in the background is the modified UTC time**. After converting, **the files of backup and rollback, and slow query logs are displayed in the frontend of the console in Beijing time, and the monitoring time is Beijing time**.

For instances with modified system time zone, if **a subsequent expansion involves data migration**, [submit a ticket](#) for assistance.

The default system time zone for instances is China Standard Time (Beijing time) .

Single-Node Cloud Disk/Dual-Node Cloud Disk Instances:

The system time zone can be modified directly on the purchase page.

For instances with modified system time zone, the actual **data stored in the background is the modified UTC time**. After converting, **the files of backup and rollback, and slow query logs are displayed in the frontend of the console in Beijing time, and the monitoring time is Beijing time**.

For instances with modified system time zone, if **a subsequent expansion involves data migration**, [submit a ticket](#) for assistance.

Prerequisites

If your instance is a dual-node local disk instance, [submit a ticket](#) before modifying the system time zone.

If you purchase a single-node cloud disk or dual-node cloud disk, you can directly modify the system time zone on purchase page with no need for application.

Use limits

Limits on instance architecture and specification

For dual-node local disk instances, [submit a ticket](#) to request a modification of the system time zone. Modifications are only supported when purchasing instances of the 90-core 720GB specification.

Single and dual-node cloud disk instances support system time zone modifications.

Limits on time zone modification

The system time zone can be modified only during instance purchase, modifications are not supported after purchase.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Create Database Instance**.
3. On the instance purchase page, set the **System Time Zone** under the **Tag** parameter.

Tag ?

Tag key Tag value x

+ Add

If the existing tags do not meet your requirements, please go to [Create Tags/Tag Values](#)

System Time Zone

China Standard Time

Only `China Standard Time` (UTC+8) is supported by default. To modify it, [submit a ticket](#) for assistance

4. For more information on how to configure other parameters on the purchase page, see [Creating TencentDB for SQL Server Instance](#). After setting all the parameters, click **Buy Now**.
5. You can query the configured system time zone under **Basic Info** on the **Instance Details** tab.

Basic Info

Instance ID

Instance Name

Instance Remarks --

Status Running

Version SQL Server 2012 Enterprise

Architecture Dual-Server High Availability

Memory 2GB

Total Capacity 10GB

Region
[Migrate across AZs](#)

AZ

Network
[Change Network](#)

System Time Zone China Standard Time

Instance Architecture Diagram

Read-only Instance Info

Comparsion table of time zones and UTC offsets

Region	Standard Time Offset	System Time Zone
Asia/Shanghai	(UTC+08:00)	China Standard Time
Asia/Singapore	(UTC+08:00)	Singapore Standard Time
America/New_York	(UTC-05:00)	Eastern Time (US & Canada)
Asia/Seoul	(UTC+09:00)	Korea Standard Time
Asia/Bangkok	(UTC+07:00)	N. Central Asia Standard Time
Asia/Tokyo	(UTC+09:00)	Tokyo Standard Time
Asia/Taipei	(UTC+08:00)	Taipei Standard Time
Asia/Jakarta	(UTC+07:00)	SE Asia Standard Time

Europe/London	(UTC)	Universal Time Coordinated
Europe/London	(UTC)	Greenwich Mean Time

Setting Instance Maintenance Information

Last updated : 2024-01-18 17:20:34

Scenario

Maintenance period is a very important concept for TencentDB for SQL Server. To ensure the stability of your TencentDB for SQL Server instance, the backend system performs maintenance operations on the instance during the maintenance period from time to time. To minimize the potential impact on your business, we recommend that you set an acceptable maintenance period for your business instance, usually during off-peak hours.

In addition, we recommend that you perform operations involving data migration during the maintenance time, such as instance specification adjustment, instance version upgrade, and instance kernel upgrade. Currently, the maintenance period is supported by primary and read-only instances.

Take the database instance specification upgrade as an example. As this operation involves data migration, after the upgrade is completed, a momentary disconnection from the database may occur. When the upgrade is initiated, the **Switch Time** can be set to **During maintenance time**, so that the instance specification will be switched during the next **maintenance time** after the instance upgrade is completed. Note that when you select **During maintenance time** for **Switch Time**, the switch will not occur immediately after the database specification upgrade is completed; instead, the sync will continue till the instance goes into the next **maintenance time** when the switch will be performed. As a result, the overall time it takes to upgrade the instance may be extended.

Note:

Before maintenance is carried out for TencentDB for SQL Server, notifications will be sent to the contacts configured in your Tencent Cloud account by SMS and email.

Instance switch is accompanied by a disconnection from the database lasting for just seconds. Make sure that your business has a reconnection mechanism.

Directions

Setting the maintenance window

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance details page.
2. In the **Maintenance Info** section on the instance details page, click **Modify**.

Maintenance Info Modify

Maintenance Period

Tue

Maintenance Time

00:00-06:00

3. In the pop-up window, select **Maintenance Window** and **Maintenance Time** as needed and click **OK**.

Modify maintenance period and time ×

Maintenance Period *

☐ Mon ☒ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat

☐ Sun

Maintenance Time *

Start Time

00:00 🕒

Duration

6 ▼ hour

OK

Cancel

Note:

Select at least one maintenance window; otherwise, the change cannot be submitted.

Performing immediate switch

If a task is configured to be switched during the maintenance window, but you need to switch it urgently under special circumstances, you can click **Switch Now** in the **Operation** column in the [TencentDB for SQL Server console](#).

Note:

Immediate switch is applicable to operations involving data migration such as instance specification adjustment, instance version upgrade, instance kernel upgrade, and cross-AZ migration.

Multi-AZ Disaster Recovery

Last updated : 2024-01-18 17:20:33

Multi-AZ Deployment

Multi-AZ deployment protects your database against database instance failures and AZ outages. For more information, see [Regions and AZs](#).

In TencentDB for SQL Server, multiple AZs are combined into a single multi-AZ to ensure high availability and failover capability of database instances.

Failover

TencentDB for SQL Server will handle failover automatically, so you can quickly restore the database operations without administrative intervention. If any of the following conditions occurs, the primary database instance will automatically switch to the replica in the replica AZ.

AZ outages.

Primary database instance failure.

Note :

No matter whether the cluster instances are deployed in multiple AZs, each TencentDB for SQL Server instance has a replica server that supports real-time hot backup to ensure the high availability of the database.

In multi-AZ deployment, TencentDB for SQL Server will automatically preset and maintain a synced replica in different AZs.

The primary database instance will be synchronously replicated across AZs to the replica to provide data redundancy, eliminate I/O freezes, and minimize latency peak during the system backup.

Purchasing a multi-AZ instance

1. Log in to the [TencentDB for SQL Server console](#), click **Create Instance** in the instance list to enter the purchase page.
2. On the TencentDB for SQL Server purchase page, select a supported region, and select a desired replica AZ in the **Multi-AZ** option.

Note :

Only certain AZs can be selected as a replica AZ. For more information, see the purchase page.

Multi-AZ

Yes

No

Network sync delay may be 2-3 ms longer if the primary and the replica reside at different AZs

Database Type	AZ
Primary	Guangzhou Zone
Replica 1	Guangzhou Zone

3. Confirm the information you enter, click **Buy Now**. After the purchase is completed, you can return to the instance list to view the newly purchased multi-AZ instance.

Upgrading to Multi-AZ

the TencentDB for SQL Server instances of Tencent Cloud Database support an upgrade from non-multi-availability zone disaster recovery to multi-availability zone disaster recovery.

Note :

Upgrading from single-AZ to multi-AZ is free of charge.

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, select the desired instance, and click **Adjust Configurations** to enter the adjustment page.

2. Select a replica AZ in the **Multi-AZ Deployment** option on the configuration adjustment page.

Note :

You can't adjust the multi-AZ configuration when using the publish/subscribe service.

As the resources are limited in multi-AZ, you may fail to purchase the resources, and you can [contact us](#) for assistance.

Upgrading to multi-AZ deployment will involve instance migration without affecting the normal use. A switch will be performed after migration is completed, which causes a momentary disconnection for few seconds.

Instance Name:

Instance Type:

Version:

Current Specs: 1-core, 2 GB MEM/10 GB disk (0 GB used)

Change Specs:

Disk Capacity: 10 GB (Increment: 10 GB)

Multi-AZ deployment:

Database Type	AZ
Primary	<input type="button" value="Guangzhou Zone 3"/>
Replica 1	<input type="button" value="Guangzhou Zone 6"/>

Time to Take Effect:

☒ During configuration adjustment, the configuration will be upgraded by migrating data. The more the data, the longer the migration. During the migration, occur, causing a flash database disconnection. Therefore, your business should have a reconnection mechanism. During the disconnection, most database, a peak hours.

Configuration Fees: USD/hour Original Price: USD/hour (Instance [Billing Details](#))

3. Click **OK**. Upgrading to a multi-AZ can be completed when the instance is purchased.

Relevant Documentation

To migrate the instance to another AZ in the same region, see [Migrating Across AZs](#).

Restarting Instance

Last updated : 2024-08-14 10:10:24

This document describes how to restart an instance in the console.

Scenario

Instance restart is a common maintenance method for TencentDB for SQL Server. It is similar to restarting a local database.

Instructions for Sensitive Operations with MFA Integrated

To enhance the security of the cloud account, TencentDB for SQL Server supports MFA (Multi-Factor Authentication), which brings an extra layer of protection in addition to the username and password. After the MFA device verification is enabled, when you perform operations such as destroying instances/changing network/modifying IP addresses/resetting passwords/deleting accounts/resetting instances/performing primary-replica switch, a second identity verification based on the MFA dynamic code will be conducted. Such operations can be executed only after successful verification. For an introduction to MFA and how to enable the operation protection, see [MFA Devices](#).

Limits

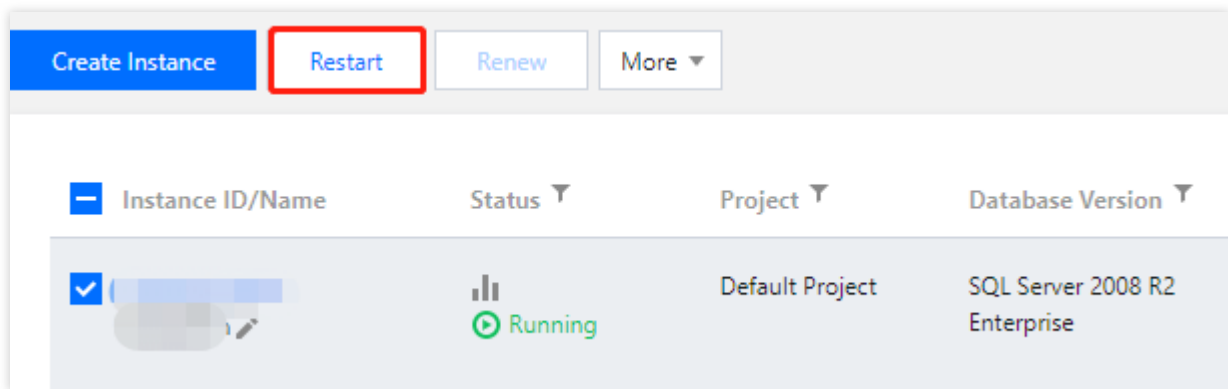
During the restart, the TencentDB for SQL Server instance cannot provide services. Therefore, make sure that an instance has stopped accepting business requests before restarting it. During the restart, if the business write volume is high, the restart may fail.

Restarting an instance does not change its physical attributes, so the private IP and any data stored in the instance will remain unchanged.

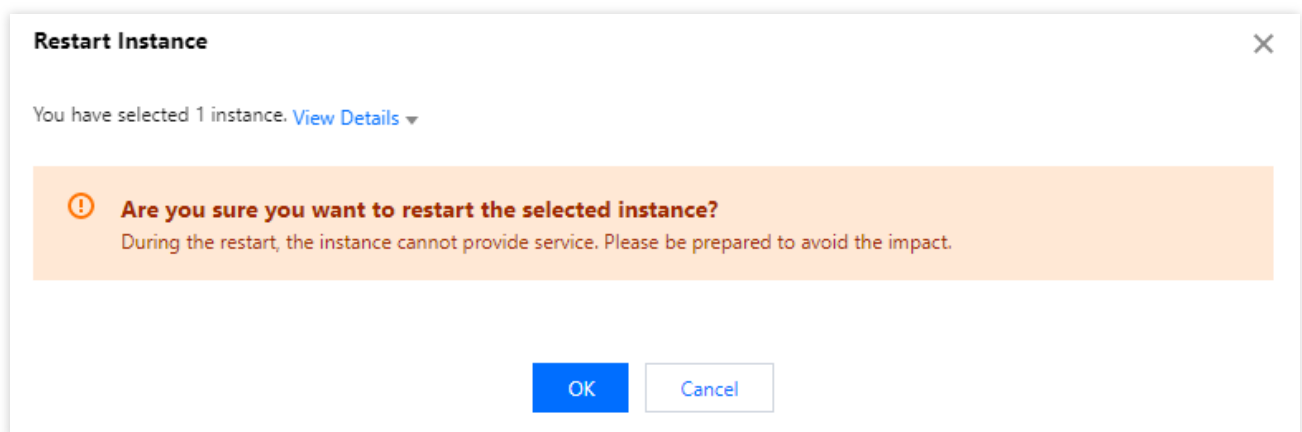
After the restart, reconnection to the database is needed. Make sure your business has a reconnection mechanism. Restart the instance during off-peak hours to ensure success and minimize the impact on your business.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select a region at the top, select one or multiple instances to be restarted, and click **Restart** at the top.



3. In the pop-up window, confirm the selected instances and click **OK**.



4. Once the instance status changes from **Restarting** to **Running**, the restart is completed.

Terminating Instance

Last updated : 2024-08-14 10:13:05

Overview

You can return pay-as-you-go instances in the console based on your business needs in a self-service manner. After a pay-as-you-go instance is returned, it will be moved to the TencentDB recycle bin and retained there for 24 hours. During the retention period, the instance cannot be accessed, but it can be restored after renewal. When an instance is returned and its status has changed to **Isolated**, it will no longer generate fees.

Note:

After the instance is terminated, its data cannot be recovered, and its backup files will also be terminated, so the data cannot be restored in the cloud. Store your backup files safely elsewhere in advance.

After the instance is terminated, its IP resources will be released simultaneously. If the instance has read-only and publish/subscribe configuration:

Read-only instances will be terminated at the same time.

After the instance is terminated, the existing pub/sub configuration on the instance will be deleted.

After the instance is terminated, the refund procedures are as detailed below:

For instances that met the 5-day no-questions-asked refund policy, the payment will be returned to your Tencent Cloud account.

For normal instances, the payment will be returned to your Tencent Cloud account by the proportion of the cash and gift cards paid for the purchase.

For orders from promotional reward channel, the refund will be charged 25% of their actual cash payment amount. These types of orders do not support self-service refunds, and you need to request a refund.

Instructions for Sensitive Operations with MFA Integrated

To enhance the security of the cloud account, TencentDB for SQL Server supports MFA (Multi-Factor Authentication), which brings an extra layer of protection in addition to the username and password. After the MFA device verification is enabled, when you perform operations such as destroying instances/changing network/modifying IP addresses/resetting passwords/deleting accounts/resetting instances/performing primary-replica switch, a second identity verification based on the MFA dynamic code will be conducted. Such operations can be executed only after successful verification. For an introduction to MFA and how to enable the operation protection, see [MFA Devices](#).

Directions

1. Log in to the [TencentDB for SQL Server Console](#), select the target instance in the instance list, and select **More > Terminate/Return** in the "Operation" column.

Instance ID / Name	Running Stat...	Project	Database Version	Configuration	Network	Availability ...	Private IP	Billing Mode	Tag
[REDACTED]	Running	Default Project	SQL Server 2016 Enterprise	1-core-2GB/20GB Basic Edition - High-Performance Cloud Disk	Default-VPC - Default-Subnet	Guangzhou Zone 4	[REDACTED]	Pay as you go	
[REDACTED]	Running	Default Project	SQL Server 2008 Enterprise	1-core-2GB/20GB Basic Edition - High-Performance Cloud Disk	Default-VPC - Default-Subnet	Guangzhou Zone 4	[REDACTED]	Pay as you go	

2. In the pop-up dialog box, indicate your consent and click **Terminate**.

Terminate Instance

1 instance is selected for termination :

ID	Attribute
[REDACTED]	Basic Edition

After the instance is completely terminated, the data will not be recovered. Please back up the instance data in advance.

After the instance is completely terminated, the IP resources are reclaimed at the same time. If the instance has associated read-only instances or pub/sub configuration:

- read-only instances will be terminated at the same time
- After the instance is terminated, the existing pub/sub configuration on the instance will be deleted.

After the instance is completely terminated, the refund procedures are as detailed below:

- The amount refunded without any reason will be refunded to the original payment account in 5 days.
- The normal self-refund amount will be returned to your Tencent Cloud account by the proportion of the cash and voucher amount paid for the purchase.
- For orders from promotional reward channel, the refund will be charged 25% of their actual cash payment amount. These types of orders do not support self-service refunds, please submit a ticket to request a refund.

☒ I have read and agreed to [Termination Rules](#)

TerminateCancel

Migrating Across AZs

Last updated : 2024-01-18 17:20:33

Scenario

You can migrate an instance to another AZ within the same region. All attributes and configurations (including the private network address and the subnet) of the instance remain unchanged after migration. The amount of time required is proportional to the volume of data in the instance. The more data there is, the longer the data migration takes. In addition, the instance access is not affected during migration.

Supported Instance Types

Single-node (formerly Basic Edition) and two-node (formerly High Availability/Cluster Edition) instances.

Note:

If a two-node (formerly High Availability/Cluster Edition) primary instance has read-only replicas or implements the pub/sub messaging paradigm, you need to [submit a ticket](#) to migrate it across AZs.

Read-only instances are not supported.

Prerequisites

The region where the instance resides must have multiple AZs.

Impact

An ongoing migration cannot be canceled.

The name, access IP, and access port of the instance remain unchanged after migration.

Data migration will occur during the instance's migration to another AZ. During the data migration, the instance can be accessed normally and your business will not be affected.

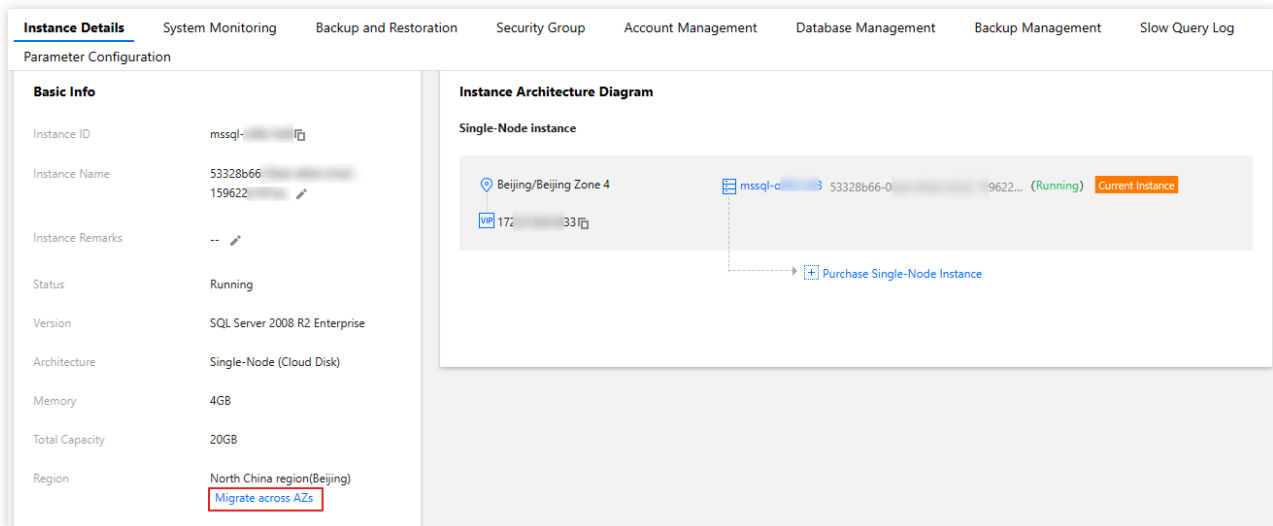
The amount of time required is proportional to the volume of data in the instance. The more data there is, the longer the data migration takes.

An instance switch will occur after migration, causing a flash disconnection. You can specify the switch time.

Directions

Migrating across AZs

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance details page.
2. In the **Basic Info** section, click **Migrate across AZs** to access the cross-AZ migration page.

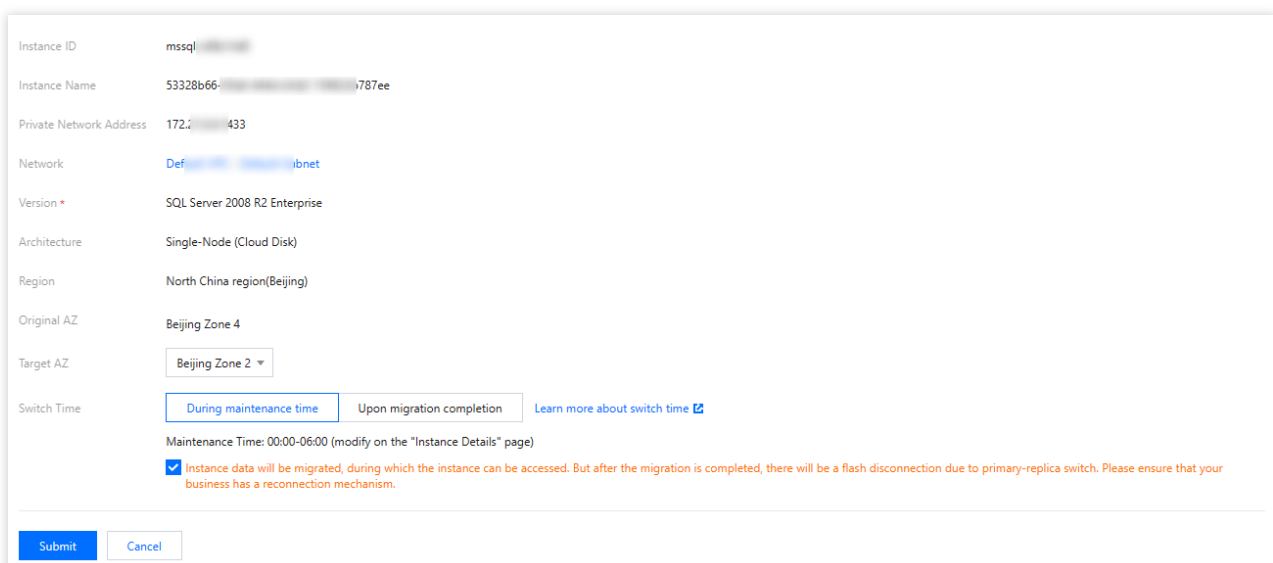


3. View the original AZ of the instance, specify the target AZ, enable/disable multi-AZ deployment, select the switch time, check the box to agree to the cross-AZ migration rules, and click **Submit**.

Note:

After you click **Submit**, instance data will be migrated to the target AZ at the underlying layer without affecting instance running and access.

After instance data is migrated, an instance switch will occur (**upon migration completion** or **during maintenance time**), causing a flash disconnection. After the switch is done, the whole process of cross-AZ migration is completed.



Viewing migration tasks

You can view cross-AZ migration tasks in the **Running Tasks** box in the upper-right corner of the **Database Management** tab.

Performing immediate switch

If the instance is scheduled to be switched during the maintenance time according to the configurations of its cross-AZ migration task, but you need to switch it urgently under special circumstances, you can click **Switch Now** in the **Operation** column in the instance list in the [TencentDB for SQL Server console](#).

Manual Primary-Secondary Switching

Last updated : 2024-08-14 10:17:38

SQL Server supports the switch between primary database and secondary database of instances. When a malfunction occurs, the system switches the secondary database to the primary to ensure system availability and data integrity. This represents the automatic switch the system performs under special circumstances such as failures. Additionally, you can manually perform the switch through the console.

Background

In enterprise-level applications, the database often forms a critical component of the business system. Any failure or downtime can have serious implications on business operations. To safeguard system availability and data integrity, high-availability solutions such as primary-secondary replication are implemented. Primary-secondary switch is a crucial technical method within the replication scheme, which enables a quick transition to the backup database when the main database encounters a problem, averting business interruptions and data losses.

Instructions for Sensitive Operations with MFA Integrated

To enhance the security of the cloud account, TencentDB for SQL Server supports MFA (Multi-Factor Authentication), which brings an extra layer of protection in addition to the username and password. After the MFA device verification is enabled, when you perform operations such as destroying instances/changing network/modifying IP addresses/resetting passwords/deleting accounts/resetting instances/performing primary-replica switch, a second identity verification based on the MFA dynamic code will be conducted. Such operations can be executed only after successful verification. For an introduction to MFA and how to enable the operation protection, see [MFA Devices](#).

Prerequisites

The instance is configured with a two-node cloud disk architecture.

Note:

If a two-node local disk architectural instance requires a manual primary-secondary switch, please [submit a ticket](#) for assistance.

The instance is in running status, with no ongoing tasks.

Points of Attention

During the primary-secondary database switch, a momentary disconnection occurs. We recommend that you run the switch during off-peak business hours and ensure your application possesses a reconnection mechanism.

After the primary-secondary database switch, the instance connection address remains the same. The original primary instance converts into a secondary instance, and the application connects automatically to the new primary instance (namely, the original secondary instance).


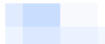
By default, the data replication mode between the primary and secondary databases is set to asynchronous replication.

Use Limits

If there are [read-only instances](#) under the primary instance or if the primary instance employs the [publish-subscribe](#) feature, then the primary-secondary switch is not supported.

The steps are as follows:

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an **instance ID** or **Manage** in the operation column to access the instance details page.

<input type="checkbox"/>	Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ
<input type="checkbox"/>	mssql- ecc3864f-cf4c-479d- b	 Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated		Guangzhou Zone 6 Primary 广州七区 Replica	:1433

2. On the instance details page, click **Primary-Replica Switch** after **Availability Info**.

Basic Info

Instance ID	mssql-
Instance Name	ecc3864 3f
Instance Remarks	--
Status	Running
Version	SQL Server 2022 Enterprise
Architecture	Two-Node (Cloud Disk)
Memory	4GB
Total Capacity	20GB
Region	South China(Guangzhou) Migrate across AZs
AZ	Guangzhou Zone 6
Network	
Private Network Address	IP: 19 Port:1433
Public Network Address	Enable

Instance Architecture Diagram

Two-Node Instance Info

Maintenance Info

[Modify](#)

Maintenance Window	Mon, Tue, Wed, Thu, Fri, Sat, Sun
Maintenance Time	00:00-06:00

3. In the Primary/Secondary switch pop-up window, select the switch time. **A momentary disconnection will occur when AZs are switched. Ensure your business possesses a reconnection mechanism** checkbox, and then click **Confirm**.

Switch Immediately: The switch will occur immediately upon the completion of data synchronization between the primary and secondary AZs.

During Maintenance Time: The switch will occur during the next maintenance window once the data synchronization between the primary and secondary AZs is complete. If the switch is scheduled for within the maintenance window, but urgent switch is needed due to unforeseen circumstances, click [Switch Now](#) in the **Operation** column for the corresponding instance in the **TencentDB for SQL Server console**. For operations related to setting the instance maintenance time, see [Setting Instance Maintenance Information](#).

Primary/Secondary Switch Log

SQL Server supports viewing primary/secondary switch logs, assisting you in understanding the detailed time and switch method, determining whether the switch is successful, and optimizing the switch strategy. For detailed operational guidance, please refer to [Querying Primary/Secondary Switch Logs](#).

Hotspot Issues

Will an instance address change after the successful switch of the primary and secondary databases of the instance?

The instance address remains the same after the switch. The original primary instance acts as a secondary one, with applications automatically connecting to the new primary instance (namely, the original secondary instance).

After a manual switch, if an upgrade or the upgrade of configurations is performed on an instance, will this affect or reset the information of the primary and secondary availability zones?

The switch will not affect or reset the information of the primary and secondary availability zones.

How can the primary and secondary databases of the primary instance be switched when read-only instances are mounted under the primary instance?

The primary instance with read-only instances does not support primary-secondary switch. If you still need to switch the primary and secondary databases, you need to release the read-only instances first, then switch the primary and secondary databases. After that, you can repurchase read-only instances.

Recycle Bin

Last updated : 2024-01-18 17:20:34

Terminated instances will be put into the recycle bin and can be restored.

Background

Tencent Cloud recycle bin offers a mechanism for repossessing cloud resources. If your account balance is sufficient, you can restore terminated instances that are still in the recycle bin.

Version description

Currently, all TencentDB for SQL Server versions support instance repossession.

Notes

Instance repossession is as described below:

For pay-as-you-go instances in the recycle bin:

Retention period: If your account has no overdue payments, terminated instances will be retained in the recycle bin for 24 hours.

Expiration processing: Instances that are not renewed before the retention period of 24 hours ends will be released and cannot be restored.

Prerequisites


The TencentDB for SQL Server instance has been terminated.

Your Tencent Cloud account balance is sufficient.

Restoring an instance from the recycle bin

1. Log in to the [TencentDB for SQL Server console](#).
2. On the left sidebar, select **SQL Server > Recycle Bin**.
3. Select the region at the top of the recycle bin page.

4. Select the target instance and click **Restore** in its **Operation** column or at the top.



5. In the pop-up window, confirm the instance information and click **OK**.

Eliminating an instance

1. Log in to the [TencentDB for SQL Server console](#).
2. On the left sidebar, select **SQL Server > Recycle Bin**.
3. Select the region at the top of the recycle bin page.
4. Find the target instance and click **Eliminate Now** in its **Operation** column.
5. In the pop-up window, confirm the instance information and click **OK**.

Note:

The instance will be completely eliminated, and its data will not be recoverable. Therefore, you need to back up the data in advance.

Adjusting Instance Configuration

Overview

Last updated : 2024-01-18 17:20:34

TencentDB for SQL Server supports flexible scaling that allows you to quickly adjust instance architecture, version, and specification in the console. You can do so at any time (at the start, during rapid development, or during peak/off-peak hours), so you can get the most out of your resources and reduce unnecessary costs in real time.

Prerequisites

You can adjust the configuration of a TencentDB for SQL Server instance and its associated instances only when they are in running status and are not executing any tasks.

Adjustment item

Adjustment Item	Description
Architecture	a. Single-node (formerly Basic Edition) instances cannot be upgraded to two-node (formerly High Availability/Cluster Edition) instances. If needed, you can purchase new two-node instances and migrate the data with DTS .
Version	a. Both single-node (formerly Basic Edition) and two-node (formerly High Availability/Cluster Edition) instances support version upgrade. b. To upgrade the version of a two-node (formerly High Availability/Cluster Edition) instance associated with a read-only instance, disassociate the read-only instance first before upgrading or submit a ticket for assistance.
Instance specification	a. All instance types support specification upgrade and downgrade. b. The upper limit of disk capacity under the corresponding specification is as displayed on the Adjust Configuration page in the console.
Disk capacity	a. Two-node (formerly High Availability/Cluster Edition) instances of local disk edition support disk capacity expansion and reduction. b. Two-node (formerly High Availability/Cluster Edition) instances of cloud disk edition support disk capacity expansion but not reduction. c. Single-node (formerly Basic Edition) instances support disk capacity expansion but not reduction.

Note:

If you need to horizontally scale the read capability of your database, use read-only instances to mitigate the pressure on the primary instance as instructed in [Managing Read-Only Instance](#).

Restrictions on upgrade/downgrade

Instance Type	Disk Type	Version Upgrade	Architecture Upgrade	Specification Upgrade	Disk Capacity Expansion	Version Downgrade	Arc Dov
Single-node (formerly Basic Edition) instance	Cloud disk	✓	X	✓	✓	X	X
Two-node (formerly High Availability/Cluster Edition) instance	Local disk	✓	✓	✓	✓	X	X
	Cloud disk	✓	✓	✓	✓	X	X
Read-only instance	Local disk	X	X	✓	✓	X	X
	Cloud disk	X	X	✓	✓	X	X

Instance disk space description

Two-node (formerly High Availability/Cluster Edition) instance - local disk

When the size of the data stored in an instance exceeds its storage space, features such as database import and rollback will become unavailable. You will need to expand its capacity or delete some database tables in the console to release the storage space. We recommend that you check the disk monitoring metrics of the instance in time as instructed in [Viewing Monitoring Charts](#) and [set the alarm policy](#) for the disk in CM, so as to prevent disk overruns from affecting your business operations.

Two-node (formerly High Availability/Cluster Edition) instance - cloud disk

When the size of the data stored in an instance exceeds its storage space, the database will become read-only. We recommend that you check the disk monitoring metrics of the instance in time as instructed in [Viewing Monitoring Charts](#), [set the alarm policy](#) for the disk in CM, and expand the storage space or delete tables in the console promptly,

so as to prevent the read-only status from affecting your business operations. The instance will become readable/writable after your expand or free up the storage space.

Single-node (formerly Basic Edition) instance - cloud disk

When the size of the data stored in an instance exceeds its storage space, the database will become read-only. We recommend that you check the disk monitoring metrics of the instance in time as instructed in [Viewing Monitoring Charts](#), [set the alarm policy](#) for the disk in CM, and expand the storage space or delete tables in the console promptly, so as to prevent the read-only status from affecting your business operations. The instance will become readable/writable after your expand or free up the storage space.

Configuration adjustment rules

You cannot cancel a configuration adjustment operation in progress.

The name, access IP, and access port of an instance will remain the same after configuration adjustment.

Data migration may be involved in configuration adjustment. During data migration, the TencentDB for SQL Server instance can be accessed normally and the business will not be affected.

Instance switchover may be needed after configuration adjustment is completed (i.e., the TencentDB for SQL Server instance may be disconnected for seconds). We recommend that you use applications configured with automatic reconnection feature and conduct the switch during the instance maintenance time. For more information, see [Setting Instance Maintenance Information](#).

Limits

As a single-node (formerly Basic Edition) instance has only one database node and no secondary node as a hot backup, when the node fails or performs tasks such as configuration adjustment and version upgrade, it will become unavailable for a long time. If your business has high requirements for the database availability, we recommend that you use two-node (formerly High Availability/Cluster Edition) instances instead.

Billing

For more information, see [Instance Adjustment Fees Description](#).

Directions

[Adjusting Instance Architecture](#)

[Adjusting Instance Version](#)

Adjusting Instance Specification

Adjusting Instance Version

Last updated : 2024-01-18 17:20:33

This document describes how to adjust the instance version in the TencentDB for SQL Server console.

Supported versions

TencentDB for SQL Server single-node (formerly basic edition) and two-node (formerly high-availability/cluster edition) instances support version upgrade.

Prerequisites

You can adjust the configuration of a TencentDB for SQL Server instance and its associated instances only when they are in running status and are not executing any tasks.

Note

Currently, the database version cannot be downgraded back after an upgrade.

Impact of configuration adjustment

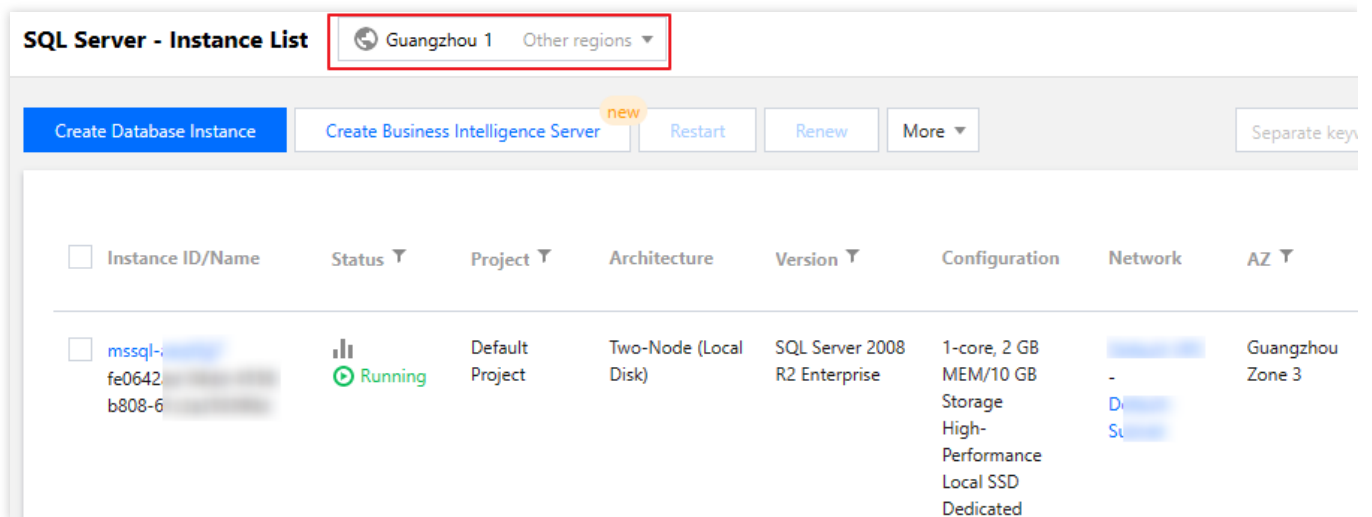
Data migration will be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.

A switch will be performed after migration is completed, which causes a momentary disconnection from the database for a few seconds. Make sure that your business has a reconnection mechanism.

Most database, account, and network operations are unavailable during the momentary disconnection. Therefore, we recommend that you perform the switch during off-peak hours.


Directions

1. Log in to the [TencentDB for SQL Server console](#), select the region and target instance in the instance list, and click **Adjust Configuration** in the **Operation** column.



SQL Server - Instance List Guangzhou 1 Other regions ▼

[Create Database Instance](#) [Create Business Intelligence Server](#) new [Restart](#) [Renew](#) [More ▼](#) [Separate keys](#)

<input type="checkbox"/>	Instance ID/Name	Status ▼	Project ▼	Architecture	Version ▼	Configuration	Network	AZ ▼
<input type="checkbox"/>	mssql-fe0642b808-6	 Running	Default Project	Two-Node (Local Disk)	SQL Server 2008 R2 Enterprise	1-core, 2 GB MEM/10 GB Storage High-Performance Local SSD Dedicated	-	Guangzhou Zone 3

2. In the **Adjust Configuration** window that pops up, select the desired **Database Version** and click **Confirm**.

Note:

Select the **Time to Take Effect** of the configuration adjustment and click the orange note.

During maintenance time: You can modify the maintenance time on the instance details page.

Adjust now: The configuration adjustment will be performed immediately.

Instance Name: mssql-

Instance Architecture: Single-Node Two-Node

Version * SQL Server 2008 R2 Enterprise SQL Server 2012 Enterprise SQL Server 2016 Enterprise SQL Server 2017 Enterprise

Current Specs: 1-core, 2 GB MEM/10 GB disk (1 GB used)

Change Specs: 1-core 2 GB MEM ▼

The maximum memory usage over the last seven days is 3 GB. For more information, click [View Details](#). The specification after being downgr

Disk Capacity: 10 500 1000 1500 2000 2500 3000 10 + GB (Increment: 10 GB)

Multi-AZ deployment: Yes No

Time to Take Effect: During maintenance time Adjust now

Maintenance Time: 00:00-06:00 (modify on the "Instance Details" page)

☒ During configuration adjustment, the configuration will be upgraded by migrating data. The more the data, the longer the migration. Du migration is completed, a switch will occur, causing a flash database disconnection. Therefore, your business should have a reconnection and network operations cannot be performed. Please switch during off-peak hours.

Configuration Fees: USD/hour Original Price: USD/hour ([Instance Billing Details](#))

OK Cancel

3. For pay-as-you-go instances, click **Adjust** in the window.

4. You will be redirected to the instance list. When the instance status changes from **Adjusting instance configuration** to **Running**, the version upgrade is completed. You can query and check the new version on the **Instance List** or **Instance Details** page.

Adjusting Instance Specification

Last updated : 2024-05-10 16:21:27

Instance Specifications

For more information on instance specifications supported by TencentDB for SQL Server and their prices, see [Product Pricing](#).

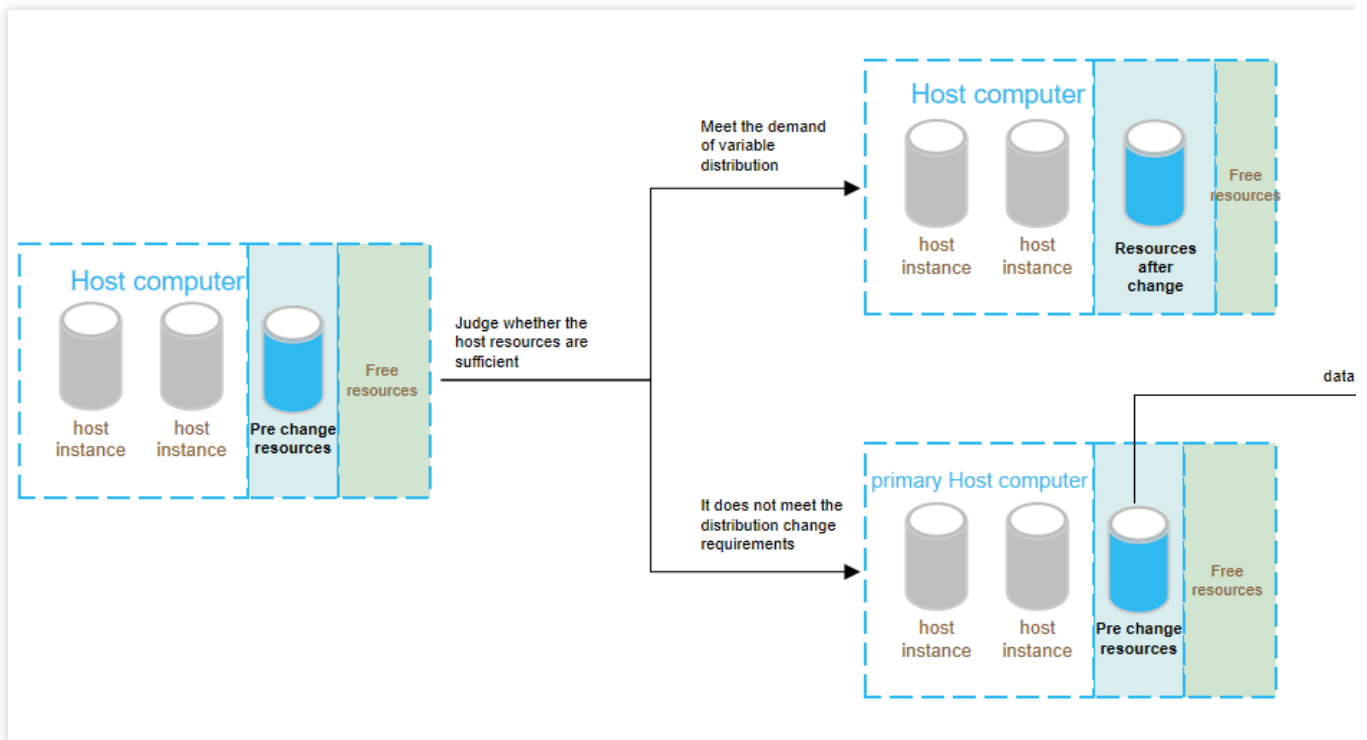
Prerequisites

You can adjust the configuration of a TencentDB for SQL Server instance and its associated instances only when they are in running status and are not executing any tasks.

Specification Adjustment Process

Scaling

After you adjust the configuration in the console, the system will determine whether to complete the adjustment by in-place scaling or data migration. The configuration adjustment process is as follows:



Impact of Configuration Adjustment

Single-item adjustment

Single-item adjustment refers to the adjustment of either the specification or disk, which has the following impact:

Instance Architecture	Disk Type	Adjustment Item	Adjustment Decision	Adjustment Impact	Time
Single-node instance (formerly basic edition)	Cloud disk	Specification upgrade or disk capacity expansion or specification downgrade	Upgrade and downgrade	1. The instance will be restarted during configuration adjustment. 2. The service will become unavailable for about 3 minutes. 3. Perform the adjustment during off-peak hours.	You need to select the time to take effect.
Two-node instance (formerly high-availability edition)	Local disk	Disk capacity reduction	Downgrade	1. The instance disk capacity will be reduced during the configuration adjustment. 2. Data migration and instance restart will not	You can select the time to take effect.

				<p>be performed, and no momentary disconnections will occur.</p> <p>3. You can custom the effective time that is also the adjustment time without affecting the business.</p>	
		Specification downgrade	Downgrade	<p>1. The instance specification will be downgraded during configuration adjustment.</p> <p>2. The service will become unavailable for around 1 minute.</p> <p>3. Perform the adjustment during off-peak hours.</p>	You need to select the time to take effect.
		Specification upgrade or disk capacity expansion	Migration and upgrade	<p>1. Data migration will be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.</p> <p>2. A switch will be performed after migration is completed, which causes a momentary disconnection from the database for few seconds. Make sure that your business has a reconnection mechanism.</p> <p>3. Most database, account, and network operations are</p>	You need to select the time to take effect.

				unavailable during the momentary disconnection. Therefore, we recommend that you perform the switch during off-peak hours.	
		Specification upgrade or disk capacity expansion	Upgrade	<p>1. The instance specification or disk capacity will be upgraded or expanded during configuration adjustment.</p> <p>2. Data migration and instance restart will not be performed, and no momentary disconnections will occur.</p> <p>3. You can customize the effective time that is also the adjustment time without affecting the business.</p>	You can select the time to take effect.
Two-node instances (formerly high-availability edition)	Cloud disk	Specification upgrade or disk capacity expansion or specification downgrade	Migration, upgrade, and downgrade	<p>1. Data migration will be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.</p> <p>2. A switch will be performed after migration is completed, which causes a momentary disconnection from the database for 30 seconds to 2 minutes. Make sure that your business has a</p>	You need to select the time to take effect.

				<p>reconnection mechanism.</p> <p>3. Most database, account, and network operations are unavailable during the momentary disconnection. Therefore, we recommend that you perform the switch during off-peak hours.</p>	
--	--	--	--	--	--

Combo adjustment

Combo adjustment refers to the adjustment of both the specification and disk capacity, which has the following impact:

Instance Architecture	Disk Type	Adjustment Item	Adjustment Decision	Adjustment Impact	Time
Single-node instance (formerly basic edition)	Cloud disk	Specification upgradeDisk capacity expansionSpecification downgrade	Upgrade and downgrade	<p>1. The instance will be restarted during configuration adjustment.</p> <p>2. The service will become unavailable for around 3 minutes.</p> <p>3. Perform the adjustment during off-peak hours.</p>	You need to select the time to take effect.
Two-node instance (formerly high-availability edition)	Local disk	Specification upgrade + disk capacity expansion	Upgrade	<p>1. The instance specification or disk capacity will be upgraded or expanded during configuration adjustment.</p> <p>2. Data migration and instance restart will not be performed, and no momentary disconnections will occur.</p> <p>3. You can custom the effective time that</p>	You can select the time to take effect.

				is also the adjustment time without affecting the business.	
		Specification upgrade + disk capacity expansion	Migration and upgrade	<p>1. Data migration will be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.</p> <p>2. A switch will be performed after migration is completed, which causes a momentary disconnection from the database for few seconds. Make sure that your business has a reconnection mechanism.</p> <p>3. Most database, account, and network operations are unavailable during the momentary disconnection. Therefore, we recommend you perform the switch during off-peak hours.</p>	You need to select the time to take effect.
		Specification upgrade + disk capacity reduction	Upgrade and downgrade	1. The instance specification will be upgraded and the disk capacity will be reduced during configuration adjustment.	The time to take effect can be customized.

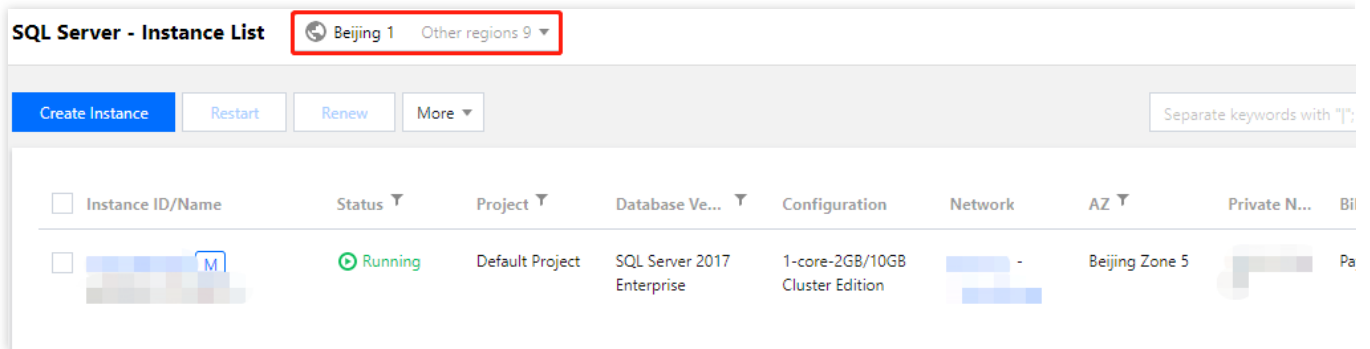
				<p>2. Data migration and instance restart will not be performed, and no momentary disconnections will occur.</p> <p>3. The time to take effect (i.e., the time when the configuration adjustment will occur) can be customized without affecting the business.</p>	
		Specification upgrade + disk capacity reduction	Migration, upgrade, and downgrade	<p>1. Data migration will be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.</p> <p>2. A switch will be performed after migration is completed, which causes a momentary disconnection from the database for few seconds. Make sure that your business has a reconnection mechanism.</p> <p>3. Most database, account, and network operations are unavailable during the momentary disconnection.</p>	You need to select the time to take effect.

				Therefore, we recommend you perform the switch during off-peak hours.	
		Specification downgrade + disk capacity expansion	Upgrade and downgrade, or migration, upgrade, and downgrade	<p>1. The service may become unavailable for around 1 minute, and data migration may be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.</p> <p>2. A switch will be performed after migration is completed, which causes a momentary disconnection from the database for few seconds. Make sure that your business has a reconnection mechanism.</p> <p>3. Most database, account, and network operations are unavailable during the momentary disconnection. Therefore, we recommend you perform the switch during off-peak hours.</p>	You need to select the time to take effect.
		Specification downgrade + disk capacity reduction	Downgrade	1. The instance specification will be downgraded during	You need to select the time to take effect.

				<p>configuration adjustment.</p> <p>2. The service will become unavailable for around 1 minute.</p> <p>3. Perform the adjustment during off-peak hours.</p>	
Two-node instance (formerly high-availability edition)	Cloud disk	Specification upgradeDisk expansionSpecification downgrade	Migration, upgrade, and downgrade	<p>1. Data migration will be involved during the instance configuration adjustment, and the larger the data volume, the longer the data migration time. Access to the instance will not be affected during this period.</p> <p>2. A switch will be performed after migration is completed, which causes a momentary disconnection from the database for 30 seconds to 2 minutes. Make sure that your business has a reconnection mechanism.</p> <p>3. Most database, account, and network operations are unavailable during the momentary disconnection. Therefore, we recommend that you perform the switch during off-peak hours.</p>	You need to select the time to take effect.

Directions

1. Log in to the [TencentDB for SQL Server console](#), select the region and target instance in the instance list, and click **Adjust Configuration** in the **Operation** column.



2. In the **Adjust Configuration** window that pops up, select the desired CPU, memory, and disk capacity, and click **OK**.

Note:

?Select the **Time to Take Effect** of the configuration adjustment and click the orange note.

During maintenance time: You can modify the maintenance time on the instance details page.

Adjust now: The configuration adjustment will be performed immediately.

If disk capacity reduction is involved, in order to avoid failures, the new disk space must be greater than or equal to 2 times the currently used disk space.

If the instance to be adjusted is associated with other read-only instances and data migration is involved, [submit a ticket](#) for assistance.

Instance Name

Instance Type * **Dual-Server High Availability**

Database Version * **SQL Server 2008 R2 Enterprise** SQL Server 2012 Enterprise SQL Server 2014 Enterprise SQL Server 2016 Enterprise SQL Server 2017 Enterprise

Current Specs 1-core, 2 GB MEM/10 GB disk (1 GB used)

Change Specs * **4-core 32 GB MEM**

Disk Capacity * **410** GB (Increment: 10 GB)

Multi-AZ deployment **Yes** **No**

Time to Take Effect **During maintenance time** Adjust Now

Maintenance Time: 00:00-06:00 (modify on the "Instance Details" page)

☒ During configuration adjustment, the configuration will be upgraded by migrating data. The more the data, the longer the migration. During the migration, the instance can still causing a flash database disconnection. Therefore, your business should have a reconnection mechanism. During the disconnection, most database, account, and network opera

Configuration Fees **USD/hour** (Instance Billing Details)

OK **Cancel**

3. In the redirected payment window for monthly subscription instances, confirm the configuration details and fees, and then click **Submit Order**. For pay-as-you-go instances, click **Adjust** in the window.

4. You will be redirected to the instance list. When the instance status changes from **Switching (after configuration adjustment)** to **Running**, the specification upgrade is completed. You can query and check the new specification on the **Instance List** or **Instance Details** page.

Read-Only Instance

Read-Only Instance Overview

Last updated : 2024-01-18 17:20:33

Overview

In scenarios where there are many read requests but only few write requests, a single instance may not be able to handle the load of read requests, which even may affect the business. To implement the auto scaling of read capabilities and mitigate the pressure on TencentDB for SQL Server, you can create one or multiple read-only instances and use them to sustain high numbers of database reads.

Unified read/write separation addresses (i.e., read and write requests are separated automatically) are not supported currently. Read-Only instances need to be accessed with separate IPs and ports.

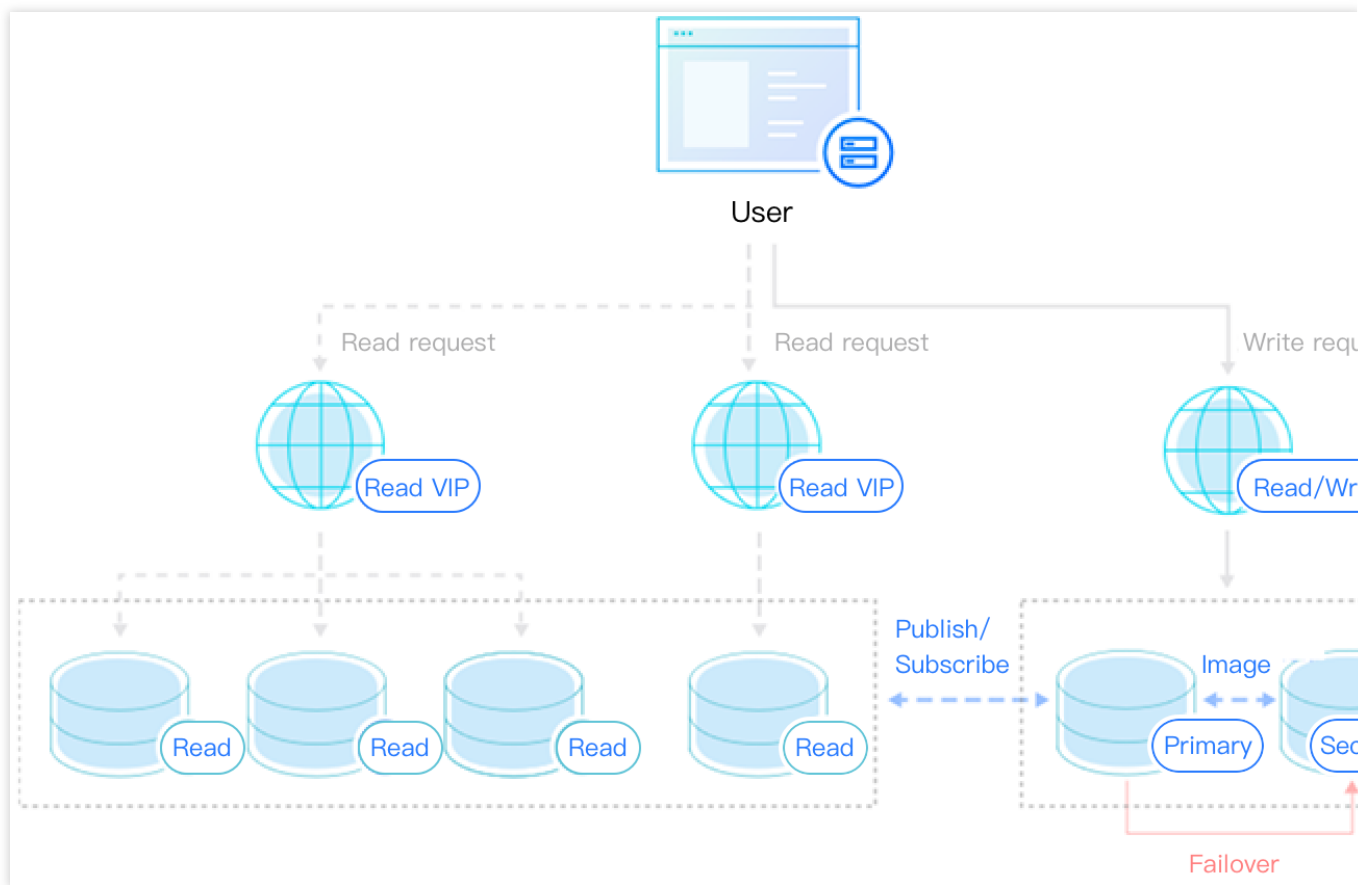
Concepts

Read-Only group: it consists of one or more load balancing-enabled read-only instances. If there are multiple read-only instances in one read-only group, read request volume can be evenly distributed among the instances. read-only groups provide IPs and ports for access to databases.

Read-Only instance: a single-node (with no replica) instance that supports read requests. It cannot exist independently; instead, it must be in a read-only group.


Architecture

Changes in the primary instance (source database) are synced to all read-only instances. Given the single-node architecture (with no replica) of read-only instances, repeated attempts to restore a failing read-only instance will be made. Therefore, we recommend you choose a read-only group rather than a read-only instance for higher availability. The read-only instance backend architecture and technology slightly vary by TencentDB for SQL Server editions: On editions below 2017 Enterprise, the publish/subscribe method is used to create read-only instances.



Note

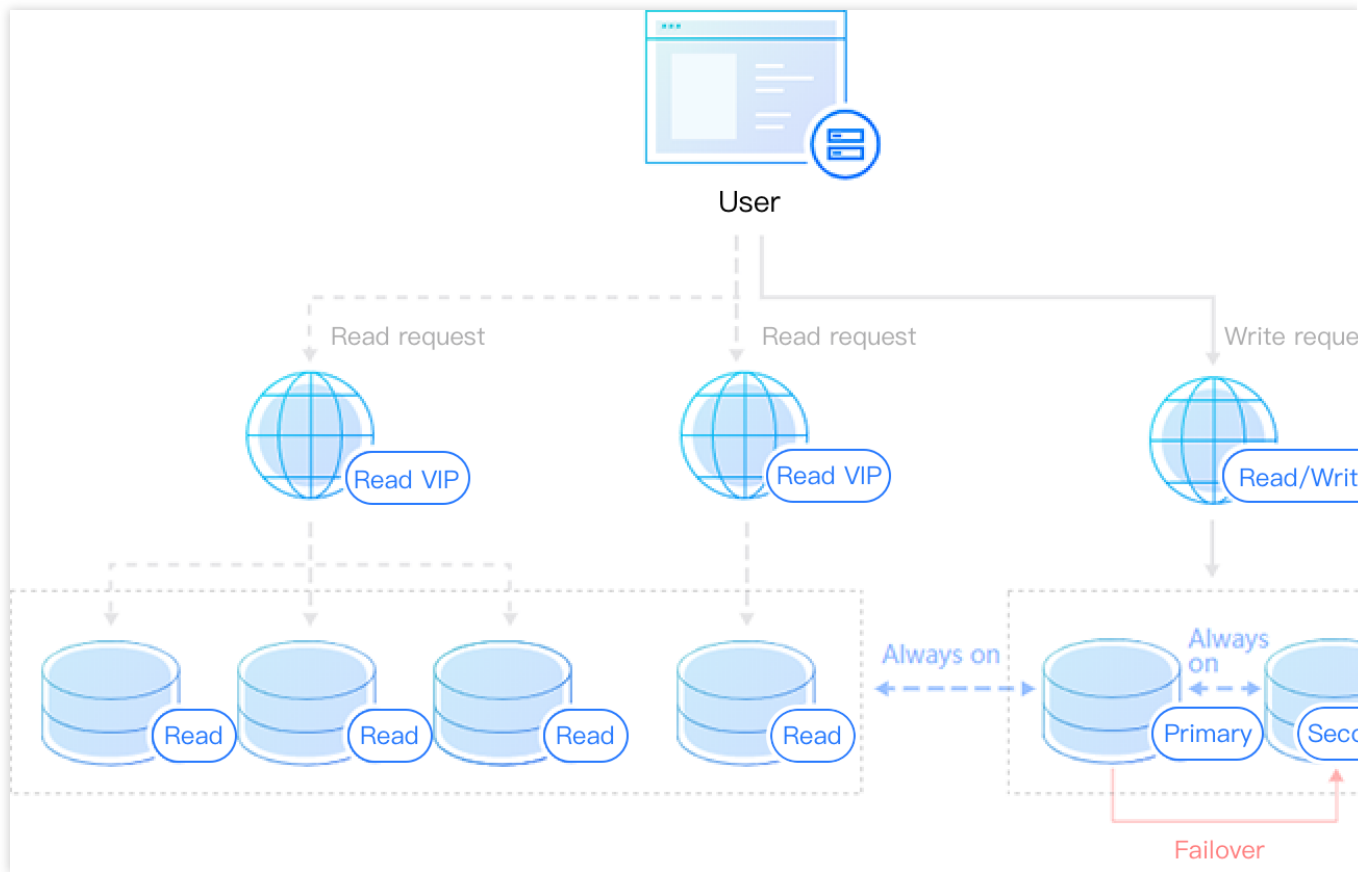
In this mode, tables without a primary key in the primary instance cannot be synced. You can use the following code to query whether there is such a table:



```
use dbname
select name from sys.sysobjects where xtype='U' and id not in(select parent_obj fro
```

If you need to create read-only instances for tables without a primary key, we recommend you use 2017 Enterprise Cluster Edition.

On 2017 Enterprise Cluster Edition and above, the Always-On method is used to create read-only instances to ensure the efficiency and stability of data sync.



Strengths

Read-Only group mode

You can connect to the VIP of a read-only group to read read-only instances in it, which can reduce the maintenance costs. You can also add the number of read-only instances in the unified read-only group to continuously expand the processing capacity of the system while ensuring the high availability of read-only instances, with no need to make any changes to the application.

Cross-AZ/region scaling

TencentDB for SQL Server supports adding read-only instances across AZs and regions, providing a low-latency, high-efficiency, and stable one-stop solution for nearby business access.

Automatic removal

The cluster management module automatically checks read-only instances. When it finds that a read-only instance is down or the delay exceeds the threshold, it stops allocating read requests to the instance and instead allocates them among the remaining healthy instances. This ensures that when a single read-only instance fails, the normal access to

the application will not be affected. When the failed instance is repaired, it will be automatically added back to the request distribution system.

Feature Limits

Up to 5 read-only instances can be created for a primary instance.

Currently, read-only instances are not supported for Standard Edition.

Currently, read-only instances cannot be added to instances in finance zones.

Read-Only instances do not support backup and rollback.

Data cannot be migrated to read-only instances.

Read-Only instances do not support database creation/deletion. If needed, please operate on a primary instance.

Read-Only instances do not support account creation/deletion/authorization and account name/password change. If needed, please operate on a primary instance.

Managing Read-Only Instance

Last updated : 2024-01-18 17:20:33

You can create, view, and terminate read-only instances in the TencentDB for SQL Server console.

Note:

Read-Only instances cannot be purchased and used separately; instead, they must be bound to a primary instance (Dual-Server High Availability Edition or Cluster Edition).

Feature Limits

Up to 5 read-only instances can be created for a primary instance.

Currently, read-only instances are not supported for Standard Edition.

Currently, read-only instances cannot be added to instances in finance zones.

Read-Only instances do not support backup and rollback.

Data cannot be migrated to read-only instances.

Read-Only instances do not support database creation/deletion. If needed, please operate on a primary instance.

Read-Only instances do not support account creation/deletion/authorization and account name/password change. If needed, please operate on a primary instance.

Creating Read-Only Instance

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the details page.
2. Click **Add Read-Only Instance** in the **Instance Architecture Diagram** on the instance details page or click **Create** on the **Read-Only Instance** page to enter the purchase page.

Basic Info

Instance ID	rmsql-2wagpazl
Status	Running
Database Version	SQL Server 2008 R2 Enterprise
Architecture	Dual-Server High Availability
Private Network Address	172.16.32.8143
Memory	2GB
Total Capacity	10GB
Region	South China (Guangzhou) Migrate across AZs

Instance Architecture Diagram

Read-only Instance Info

Guangzhou/Guangzhou Zone 7

VIP 172.16.32.8143

rmsql-2wagpazl

Master-Slave Sync

+ Add Read-only Instance

3. On the purchase page, select the desired read-only instance configuration, confirm that everything is correct, and click **Buy Now**.

Note:

If you need to unify the expiration time of the primary and read-only instances, you can set the collective expiration date in the [Renewal Management console](#).

Primary Instance Info

Instance Name	2e193fe7-3129-4074-be11-8b3516ee09fd	Instance ID	rmysql-fgs7e4c	Project	DEFAULT
Network	/	Region	Beijing	AZ	Beijing
Architecture	Cluster Edition	Instance Specs	1-core, 2 GB/10 GB	Database Version	SQL S

Specify RO Group

Assigned by system

[Learn about RO Group](#)

System automatically assigns RO groups which are not enabled with the [read-only instance removal based on the delay threshold](#). Please retain a

Billing Mode

Pay as You Go

Region

Guangzhou

Shanghai

Nanjing

Beijing

Hong Kong (China)

[Detailed Comparison](#)

The cloud service products in different regions cannot communicate with each other through private network. Thus, it is recommended to select the region closest to your customers to minimize the access latency.

AZ

Guangzhou Zone 2

Guangzhou Zone 3

Guangzhou Zone 4

Guangzhou Zone 6

Network

VPC Both RO group and source instance must in the same VPC

test

xx | 10.0.0.0/24

253 subnet IPs in total, with 252 available

You can also [Create VPCs](#) or [Create Subnets](#). After the TencentDB instance is purchased, the VPC and subnet can be modified in the console. In the current network environment, only devices in "VPC test " can access the database instance.

In the Guangzhou region and under VPC test , the servers that can be connected via private network are: 0 pcs.

Total Fees

USD/hour (Billing Details)

Buy Now

Viewing Read-Only Instance

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, instances with an **R** flag are read-only instances. Click an instance ID or **Manage** in the **Operation** column to enter the read-only instance details page.
2. In the **Instance Architecture Diagram** on the instance details page, you can view the information of the bound primary instance. You can click the instance ID to enter the details page of the primary instance. You can also enter the details page of the read-only instance from the **Instance Architecture Diagram** of the primary instance.

Note:

Some features on the read-only instance details page cannot be modified and are synced from the primary instance. If you need to change them, please do so on the primary instance details page.

Basic Info

Instance ID	mssql-figa7x4c
Status	The instance is creating RO replica
Database Version	SQL Server 2017 Enterprise
Architecture	Cluster Edition
Private Network Address	10.0.1.15/143
Memory	2GB
Total Capacity	10GB
Region	North China (Beijing) Migrate across AZs
AZ	Beijing Zone 5
Network	test-vpc-leondzhang Change Network
Tag	--
Project	DEFAULT PROJECT

Instance Architecture Diagram

Read-only Instance Info

The diagram illustrates the architecture of the read-only instance. It shows a primary instance 'mssql-figa7x4c' in Beijing/Beijing Zone 5, which is the 'Current Instance'. A read-only instance 'mssqlrg-adxfr0...' is shown in Guangzhou/Guangzhou Zone 6, with a delay of '-s'. The diagram also includes a 'Master-Slave Sync' section and options to 'Add Reac' (Add Read-only Instance).

Terminating Read-Only Instance

An read-only instance can be terminated in the same way as a primary instance as instructed in [Terminating Instance](#).

Read-Only Group

Last updated : 2024-07-30 16:21:58

Overview

TencentDB for SQL Server allows you to create one or more read-only instances to form a read-only group, which is suitable for read/write separation and one-primary-multiple-replica application scenarios and capable of greatly enhancing the read load capacity of your database.

Note :

In scenarios where an RO group contains only a single read-only instance, there exists a risk of a single point of failure. Moreover, such an RO group will not be included in the overall availability calculations for the Tencent Cloud SQL Server service. A single read-only instance does not offer an availability SLA guarantee. It is recommended to purchase at least two read-only instances for an RO group to ensure its availability.

Prerequisites

A primary instance must be created first before a read-only instance can be created. For more information, please see [Creating TencentDB for SQL Server Instance](#).

Directions

Creating Read-Only group

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the details page.
2. Click **Add Read-Only Instance** in the Instance Architecture Diagram on the instance details page to enter the purchase page.

Basic Info

Instance ID	ms-sql-2008r2-zl
Status	Running
Database Version	SQL Server 2008 R2 Enterprise
Architecture	Dual-Server High Availability
Private Network Address	172.16.32.8/1433
Memory	2GB
Total Capacity	10GB
Region	South China (Guangzhou) Migrate across AZs

Instance Architecture Diagram

Read-only Instance Info Read-only Instz

3. On the purchase page, select the desired read-only instance configuration, confirm that everything is correct, and click **Buy Now**.

Note:

If the Specify RO Group option is configured as Create RO group, the following basic information of the new read-only group should be entered on the purchase page.

RO Group Name: the group name doesn't have to be unique and can contain up to 60 letters, digits, hyphens, underscores, and dots.

Remove Delayed RO Instances: this option indicates whether to enable the removal policy. If a read-only instance's delay exceeds the threshold, it will be removed and become inactive, and its weight will be set to 0. It will be put back into the read-only group when its delay falls below the threshold. No matter whether delayed read-only instance removal is enabled, a read-only instance that is removed due to instance failure will rejoin the read-only group when it is repaired.

Delay Threshold: this sets a delay threshold for the read-only instance. When the threshold is exceeded, the instance will be removed from the read-only group.

Least RO Instances: this is the minimum number of instances that should be retained in the read-only group. When there are fewer instances in the read-only group, even if an instance exceeds the delay threshold, it will not be removed.

Specify RO Group

Create RO Group

[Learn about RO Group](#)

RO Group Name

Remove Delayed RO Instances
☒ Enable (Note that this setting only applies to delayed RO instances. Failed RO instances are always removed directly and added backed after they're recovered.) [Learn about RO Group](#)

Delay Threshold

—

10

+

s

Least RO Instances

1 pcs

When the number of purchased read-only instances is below the set number, no read-only instance is removed.

Specify Read-Only Group	Description

Assigned by system (not specified)	If multiple instances are purchased at a time, each of them will be assigned to an independent read-only group, and their weights will be automatically assigned by the system by default.
Create read-only group	Create a read-only group. If multiple instances are purchased at a time, all of them will be assigned to this new read-only group, and their weights will be allocated by the system automatically by default.
Existing read-only group	Specify an existing read-only group. If multiple instances are purchased at a time, all of them will be assigned to this read-only group. Their weights will be allocated as configured in the read-only group. If assignment by the system is set for the read-only group, the instances will be added to the group automatically according to the purchased specifications. If custom allocation is set, their weights will be zero by default. As the same private IP is shared within a read-only group, if a VPC is used, the same security group settings will be shared. If a read-only group is specified, it is not possible to customize any security group when instances are purchased.

4. Return to the instance list. The status of the created read-only instance is **Delivering**. If the status changes to **Running**, the instance has been successfully created.

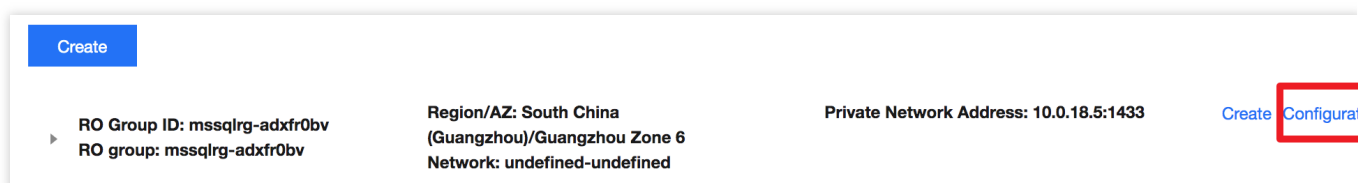
Configuring read-only group

On the read-only group configuration page, you can configure the basic information of the group such as name, removal policy, delay threshold, least read-only instances, and read weight.

Note:

Read-Only instances in a read-only group can use different specifications, and their read traffic weights can be set. Read-Only instances in the same read-only group can have different expiration dates and billing modes.

1. In the [instance list](#), select a primary instance for which to set a read-only group, and click the **Instance ID** or **Manage** in the **Operation** column on the right to enter the instance management page.
2. On the instance management page, click the **Read-Only Instance** tab and click **Configuration** to enter the read-only group configuration page.



3. In the pop-up window, configure the read-only group options.

RO group configuration

×

RO Group ID

RO Group Name

Remove Delayed RO Replicas

Note that this setting only applies to delayed RO instances. Failed RO instances are always removed directly and added backed after they're recovered.

Assign Read Weight

Assigned by system

Custom

Read-Only Instance...

Weight (an integer b...

Load Rebalancing

If load rebalancing is disabled, modifying weight only takes effect for new loads and will not affect the read-only instances accessed by the original persistent connection and not cause flash disconnection of database.

OK

Cancel

RO Group Name: the group name doesn't have to be unique and can contain up to 60 letters, digits, hyphens, underscores, and dots.

Remove Delayed RO Instances: this option indicates whether to enable the removal policy. If a read-only instance is removed, its weight will be automatically set to 0.

Delay Threshold: this sets a delay threshold for the read-only instance. When the threshold is exceeded, the instance will be removed from the read-only group.

Least RO Instances: this is the minimum number of instances that should be retained in the read-only group. When there are fewer instances in the read-only group, even if an instance exceeds the delay threshold, it will not be removed.

Assign Read Weight: The RO group supports two methods of weight setting: automatic weight assignment by the system and custom weight assignment. The read-only address of the replica node does not participate in weight assignment. The range of the input weight is 0 - 100, and the input weight must be an integer. The system automatically sets the list of read weight values for SQL Server instances as follows:

Instance Specification	Weight
2,000 MB memory	1
4,000 MB memory	2
8,000 MB memory	2
12,000 MB memory	4

16,000 MB memory	4
24,000 MB memory	8
32,000 MB memory	8
48,000 MB memory	10
64,000 MB memory	12
96,000 MB memory	14
128,000 MB memory	16
244,000 MB memory	26
488,000 MB memory	50

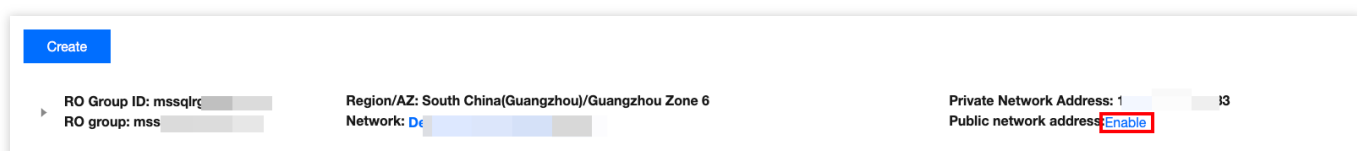
Rebalance:

Modifying weight will only affect new loads if rebalancing is disabled. The operation has no impact on read-only instances accessed by existing persistent connections and does not cause momentary database disconnection.

If rebalancing is enabled, all connections to the database will be temporarily disconnected, and the loads of newly added connections will be balanced according to the set weights.

Enabling/Disabling Public Network Addresses for Read-Only Groups

1. Log in to the [SQL Server Console](#).
2. Select the region, In the instances list, locate the primary instance for which you want to enable the public network address of the read-only group, and then click **Instance ID** or the **Manage** option in the **Operation** column.
3. On the **Instance Details** page, select the **Read-Only Instances** tab, and then click **Enable** or **Disable** following the **Public Network Address** under the **RO Group**.



4. Read the prompts displayed in the enabling or disabling public network settings window, and click **Confirm**. For more information on this operation, see [Enabling or Disabling Public Network Addresses](#).

Terminating and deleting read-only group

Read-Only groups cannot be deleted manually.

A read-only group will be automatically deleted when the last read-only instance in it is eliminated.

Empty read-only groups cannot be retained.

Backup node read-only

Last updated : 2024-04-10 17:19:31

TencentDB for SQL Server 2017/2019/2022 instances with Always On Dual-Node Cloud Disk architecture support the replica node read-only feature. Enabling the replica node read-only allows access to the replica node through its unique read-only address reserved for the replica node, sharing the read requests of the primary node. This can effectively save you the cost of an additional read-only instance. You can enable or disable the replica node read-only feature through the console, as well as modify the network settings of the read-only address for the replica node.

Prerequisites


The architecture and version of the primary instance are: TencentDB for SQL Server 2017/2019/2022 Always On Dual-Node.

The storage type of the primary instance is: Enhanced SSD / Balanced SSD.

The primary instance is running.

Note:

You can view whether the above information meets the requirements in the instance list.

<input type="checkbox"/>	Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration
Search for "Instance"					
<input type="checkbox"/>	mssql-1 696da1 9077-6	 Running	Two-Node (Cloud Disk)	SQL Server 2019 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated

The primary and replica instances must be synchronized properly.

Note:

If the instance is new, it is necessary to wait until the AG group of both nodes is established and the primary node and replica node are in normal synchronization state before you enable the replica node read-only feature.

Feature Limits

Enabling replica node read-only will be affected by the limit on the number of databases a single instance can support. If the limit is exceeded, this feature cannot be enabled. For the specific limit, please see [Maximum number of databases created in a single instance](#).

The enabling of replica node read-only feature is affected by the data synchronization between the primary and replica instances. When the data is not synchronized, the feature cannot be enabled.

The read-only mode of the replica node operates as an independent instance state and can only be accessed through its unique read-only address reserved for the replica node. It is unrelated to RO instances and the RO groups they are in; therefore, the replica node cannot join an RO group to participate in weight setting.

After replica node read-only is enabled, it supports access only via private network addresses, and enabling public network access on the replica node is not supported.

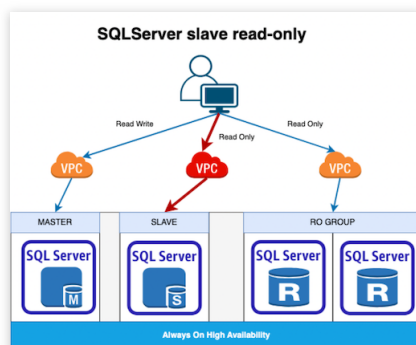
After replica node read-only is enabled, the replica node will not be removed if the replica node crashes or experiences a delay timeout. Accessing the replica node through the read-only address may result in data delays at this time. Once the node recovers, data synchronization will return to normal state.

After replica node read-only is enabled, if the instance undergoes a migration or switch across availability zones, the read-only replica node may experience a brief disconnection. Please ensure that your business has the reconnection mechanism.

After replica node read-only is enabled, if the instance undergoes a migration or configuration switch, the read-only replica node may experience a brief disconnection. Please ensure that your business has the reconnection mechanism.

After replica node read-only is enabled, if the instance undergoes HA switch, the read-only replica node may experience a brief disconnection or prolonged interrupts. Please ensure that your business has the reconnection mechanism.

Replica Node Read-Only Architecture Diagram



Enabling Replica Node Read-only

1. Log in to the [SQL Server Console](#). Under the instance list, click the **Instance ID** or the **Manage** option in the **Operation** column to enter the instance details page.
2. On the instance details page, enable **Replica Node Read-only**.

Replica Node Read-OnlyStatus ☒ Disabled

3. On the page where replica node read-only is enabled, complete the following configuration and click **OK**.

Replica Node Read-Only Enabled

1. If the read-only function of the replica node is enabled, you can access it via the read-only address of the replica node, effectively helping you save one read-only instance. For details, see [Replica Node Read-Only](#).
2. If the read-only function of the replica node is enabled, the replica node will not be removed in case of crashes or timeout due to delay. Data delay may occur if the replica node is accessed via the read-only address. When the node recovers, data synchronization will return to normal.
3. If the read-only function of the replica node is enabled, the number of databases a single instance can host will be limited. For details, see [Database Quantity Limit](#).
4. If the read-only function of the replica node is enabled, the replica node will be momentarily disconnected when instances are migrated or switched across availability zones or with configuration adjusted. During the HA switch of the instance, the replica node may experience momentary or long-term disconnection. Please ensure that your services have a reconnection mechanism.

Select Network

De De CIDR: 1 0, subnet IP/available IP: 4093/4055In the current network environment, only CVM in the "I " can access this database [Create VPC](#)
[Create Subnet](#)

Read-Only Address

☒ Auto-Assign IP ☐ Specify IP

OK

Cancel

Parameter	Description
Select a network	Select VPC network. The VPC, Subnet, and Read/Write address are the same as those of the primary node by default. You can also manually select them. After confirming the VPC, only hosts corresponding to the VPC can access the database.
Read-only Address	Auto-Assign IP: The system automatically assigns an IP address. Specify IP: You can manually define the Subnet IP address.

Note:

The replica node read-only feature is enabled for the primary node. When the replica node switches, the read-only address will be bound to the new replica node.

4. Once the instance status returns to **Running**, the read-only setting for the replica node will be completed.

Changing the Replica Node Read-only Network

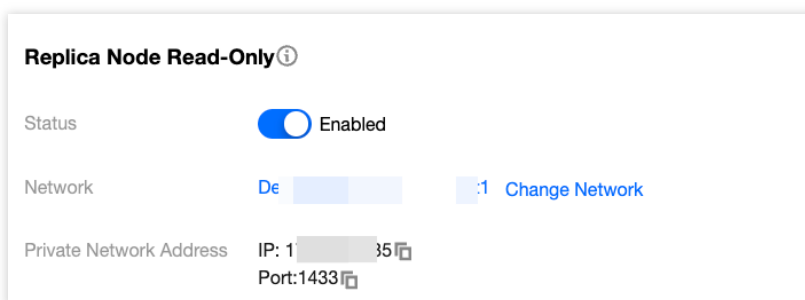
After the replica node read-only is enabled, if you want to change the network of the replica node read-only address, please see [Changing Network](#).

Disabling Replica Node Read-only

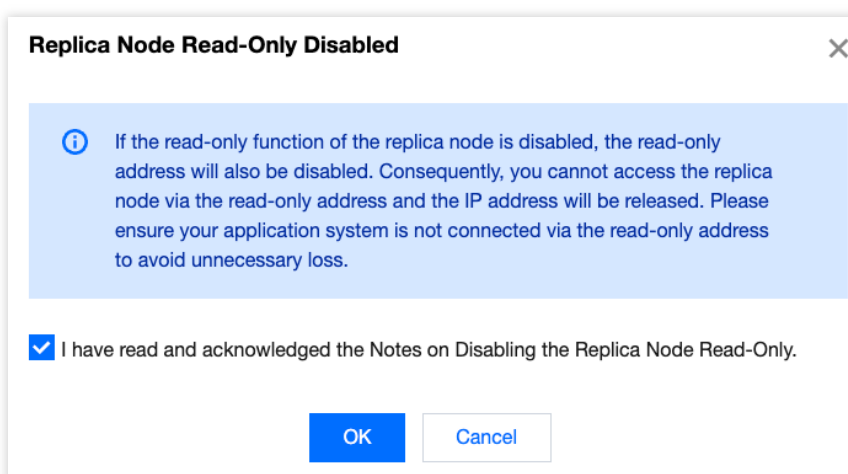
Note:

After the replica node read-only feature is disabled, the read-only address will be closed, and you will not be able to access the replica node via the read-only address. Additionally, the IP address will be released. If you want to disable the replica node read-only feature, please ensure your application system does not access the network through the read-only address to avoid losses.

1. Log in to the [SQL Server console](#). In the instances list, click the **Instance ID** or the **Manage** option in the **Operations** column to enter the instance details page.
2. On the instance details page, disable **Replica Node Read-Only**.



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



FAQs

After the replica node read-only feature is enabled, what impact will be on the application system's access to the replica node read-only address if the primary node or the replica node fails?

After the replica node read-only feature is enabled, **if the primary node fails and triggers a Primary-Replica Switch**, access via the read-only address will not be provided until the original primary node recovers to prevent the original replica node from being overwhelmed by a sudden surge in traffic. Consequently, the corresponding application system temporarily cannot access the data of the replica node, until the original primary node is recovered and switched to the new replica node, read-only access to the replica node can be resumed;

After the replica node read-only feature is enabled, **if the replica node fails**, the replica node will not provide services temporarily, and the application system cannot access the instance through the replica node read-only address until the replica node is recovered.

Network and Security

Switching from Classic Network to VPC

Last updated : 2024-01-18 17:20:33

This document describes how to switch the instance network from classic network to VPC in the TencentDB for SQL Server console.

Network types

There are two types of TencentDB network environments: classic network and VPC.

Classic network: It is the public network resource pool for all Tencent Cloud users. All your Tencent Cloud resources will be centrally managed by Tencent Cloud.

VPC: It is a logically isolated network space that can be customized in Tencent Cloud. Even in the same region, different VPCs cannot communicate with each other by default. Similar to the traditional network in an IDC, a VPC is where your Tencent Cloud service resources are managed.

Overview

Tencent Cloud supports classic network and VPC, which are capable of offering a diversity of smooth services. On this basis, we provide the classic network to VPC switch feature to help you manage network connectivity with ease.

Supported instance types

Primary and read-only instances on all versions.

Notes

Currently, you can only switch from classic network to VPC but not vice versa.

After the switch, VPC access will take effect immediately. The original classic network access will be retained for 24 hours; therefore, other instances associated to the instance should be migrated to VPC within 24 hours so as to guarantee uninterrupted access.

If **Valid Hours of Old IP** is set to 0, the IP will be released immediately after the network is changed. Then, you can only access the instance over the VPC.


After you switch the primary instance's network, the networks of read-only instances associated with the primary instance won't be automatically switched, that is, you need to manually switch them.

Subnet description

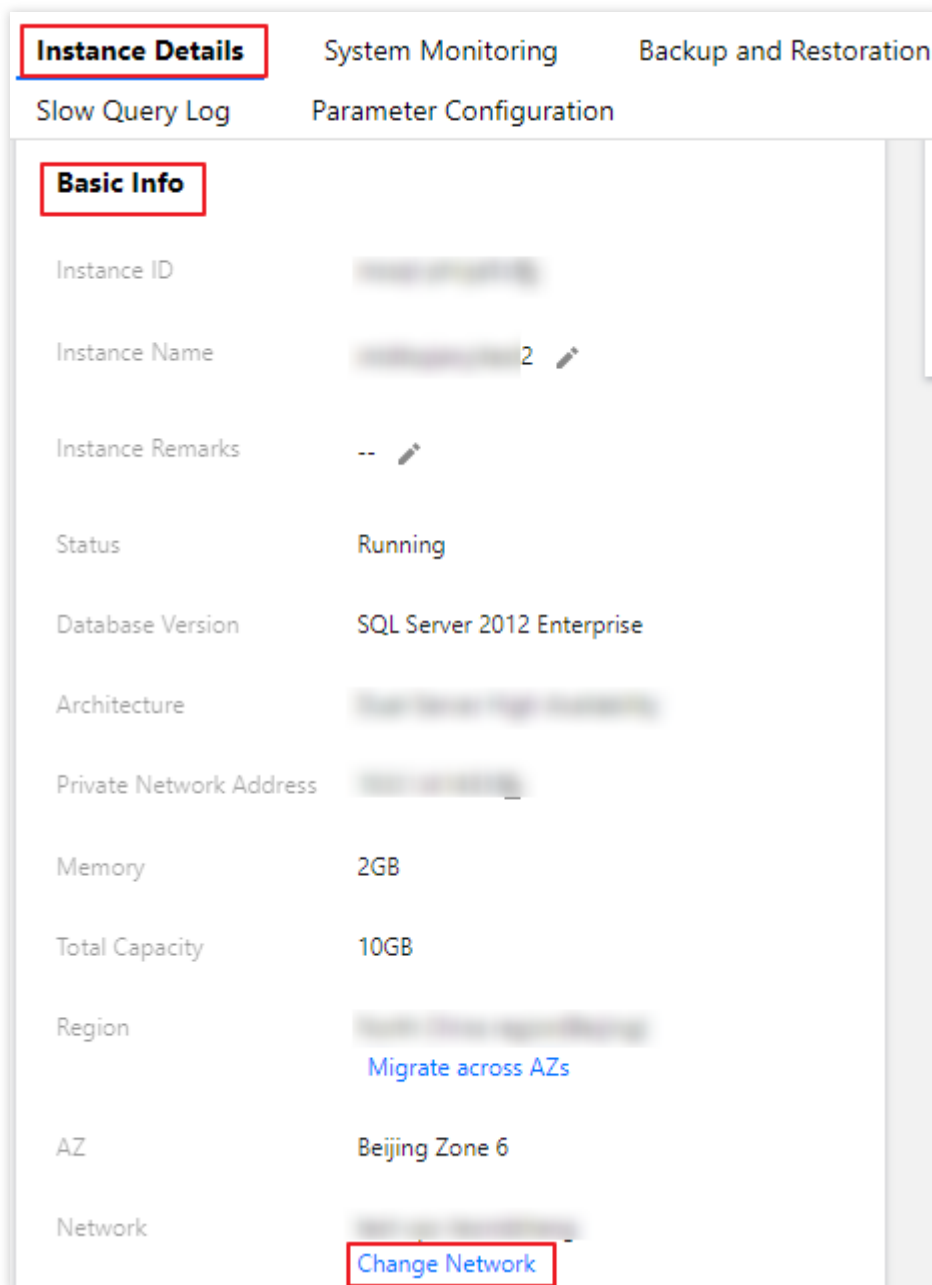
A subnet is a logical network space in a VPC. You can create subnets in different AZs in the same VPC, which communicate with each other over the private network by default. Even if you select a subnet in another AZ in the same region, the network latency will not be increased because the actual business connection adopts nearby access.

Directions

1. Log in to the [TencentDB for SQL Server console](#), select the region, and click **Instance ID/Name** of the target instance.



2. On the **Instance Details** page, click **Switch to VPC** in **Basic Info > Network**.



3. In the pop-up window, select a VPC, set **Valid Hours of Old IP**, select **Auto-Assign IP** or **Specify IP**, and click **OK**.


Select Network: After a VPC is selected, only servers in it can access the database.

Valid Hours of Old IP: You can select 0–168 hours. If it is set to 0, the old IP will be repossessed immediately.

Auto-Assign IP: The system automatically assigns the new IP address.

Specify IP: You can customize the subnet IP address.

Change Network



1. When the network is changed, the IP of the instance will be replaced with a new one. The old one will become invalid 24 hours following the change by default. Please modify the client program accordingly in time.

2. If "Valid Hours of Old IP" is set to 0 hours, the IP is released immediately after the network is changed.

3. You can only select a new VPC and subnet in the same region where the instance resides.

Select Network

CIDR: , Subnet IPs/Available IPs: 253/246

In the current network environment, only CVMs in the can access this database.[Create VPCs](#) [Create Subnets](#)

Valid Hours of Old IP

-

24

+

hour(s) after the change (Value range: 0 to 168)

☒ Auto-Assign IP

☐ Specify IP

OK

Cancel

4. The classic network is successfully switched to VPC when the instance status becomes **Running**.

Changing Network

Last updated : 2024-08-14 10:37:12

TencentDB for SQL Server supports changing the instance's network according to business needs. This document describes related operations in the SQL Server console to change the network.

Note:

TencentDB for SQL Server supports changing only the private network address with the network unchanged. For more information, see [Changing Only Private Network Addresses with Network Unchanged](#).

TencentDB for SQL Server supports manually releasing the old IP address when the old-address recycling time is set to non-zero during the private IP change. For more information, see [Manually Releasing Reserved Addresses](#).

There are two types of TencentDB network environments: VPC and classic network.

Classic network: It is the public network resource pool for all Tencent Cloud users. All your Tencent Cloud resources will be centrally managed by Tencent Cloud.

VPC: It is a logically isolated network space that can be customized in Tencent Cloud. Even in the same region, different VPCs cannot communicate with each other by default. Similar to the traditional network in an IDC, a VPC is where your Tencent Cloud service resources are managed.

Scenario

Tencent Cloud supports VPC to offer a diversity of smooth services. On this basis, we provide the VPC change feature to help you manage network connectivity with ease.

Scenario 1: VPC to VPC: Changing the instance's network; the private network address will change, and the application system needs to reconnect using the new private network address.

Scenario 2: Only changing the private network address with network unchanged: The instance's network remains unchanged, only the private network address changes, and the application system needs to reconnect using the new private network address.

Scenario 3: Manually releasing reserved addresses: When changing the internal network address and the old address recycling time is not set to 0, you can quickly release the old address and enable the new internal network address by releasing the reserved address.

Instructions for Sensitive Operations with MFA Integrated

To enhance the security of the cloud account, TencentDB for SQL Server supports MFA (Multi-Factor Authentication), which brings an extra layer of protection in addition to the username and password. After the MFA device verification is enabled, when you perform operations such as destroying instances/changing network/modifying IP

addresses/resetting passwords/deleting accounts/resetting instances/performing primary-replica switch, a second identity verification based on the MFA dynamic code will be conducted. Such operations can be executed only after successful verification. For an introduction to MFA and how to enable the operation protection, see [MFA Devices](#).

Supported instance types

VPC to VPC: primary instances (including [SSIS instances](#)), read-only instances, and replica nodes (replica node read-only addresses).

Only changing the private network address with network unchanged: primary instances (including [SSIS instances](#)) and read-only instances.

Manually releasing reserved addresses: primary instances (including [SSIS instances](#)) and read-only instances.

Note

Changing the private network IP will affect the database services currently being accessed.

Changing the network will change the instance IP. The original IP will become invalid after 24 hours by default. Modify the instance IP on the client promptly.

If **Valid Hours of Old IP** is set to 0, the IP will be released immediately after the network is changed.

Subnet description

A subnet is a logical network space in a VPC. You can create subnets in different AZs in the same VPC, which communicate with each other over the private network by default. Even if you select a subnet in another AZ in the same region, the network latency will not be increased because the actual business connection adopts nearby access.

Scenario 1 Operation Steps: VPC to VPC

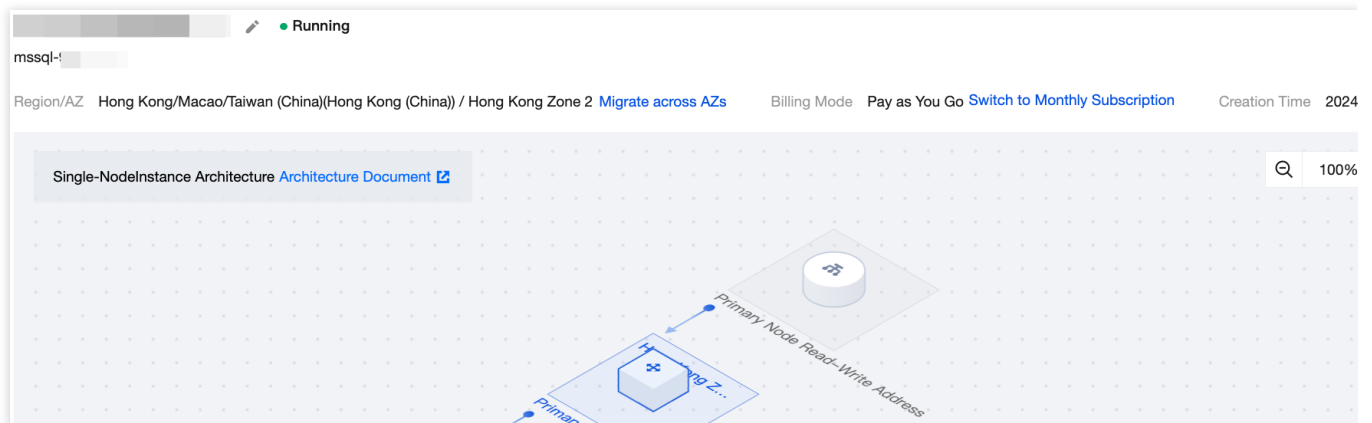
Primary Instance/Read-Only Instance

Replica Node Read-Only Address

1. Log in to the [TencentDB for SQL Server console](#), select the region, and click **Instance ID/Name** of the target instance or **Manage** in the **Operation** column.

Instance ID/Name	Status	Architecture	Version	Configuration	Network	AZ	Private Network Address	Billing
Search for "Instance ID:mssql-22shij42", 1 result found Back to list								
mssql-2	Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/30 GB Storage Enhanced SSD Dedicated		Shanghai Zone 4 Primary Shanghai Zone 2 Replica	:1433	Monthly 2024-04

2. In the **Instance Info** on the **Instance Details** page on the right-hand side click **Change Network** under the primary node or the current read-only instance.



3. In the pop-up window, select a VPC, set **Valid Hours of Old IP**, select **Auto-Assign IP** or **Specify IP**, and click **OK**.

Select Network: After a VPC is selected, only servers in it can access the database.

Valid Hours of Old IP: You can select 0–168 hours. If it is set to 0, the old IP will be repossessed immediately.

Auto-Assign IP: The system automatically assigns the new IP address.

Specify IP: You can customize the subnet IP address.

Change Network

1. When the network is changed, the IP of the instance will be replaced with a new one. The old one will become invalid 24 hours following the change by default. Please modify the client program accordingly in time.

2. If "Valid Hours of Old IP" is set to 0 hours, the IP is released immediately after the network is changed.

3. You can only select a new VPC and subnet in the same region where the instance resides.

Select Network

CIDR: 10.0.0.0/24, Subnet IPs/Available IPs: 10.0.0.1-10.0.0.254
In the current network environment, only CVMs in the "t-123456789" can access this database.[Create VPCs](#) [Create Subnets](#)

Valid Hours of Old IP

-

24

+

 hour(s) after the change (Value range: 0 to 168)

☐ Auto-Assign IP

☒ Specify IP

OK

Cancel

4. The VPC is changed successfully when the instance status becomes **Running**.

Note:

The modification of the replica node read-only address is supported only after the replica node read-only feature has been enabled. For more information on enabling the feature, see [Replica Node Read-Only](#).

Changing the network will change the replica node read-only IP, and the original IP address will be immediately repossessed.

To change the network, you can only select the VPC network and subnet in the same region where the instance resides.

1. Log in to the [SQL Server Console](#), select the region, and click the **Instance ID / Name** or **Manage** in the Operations column of the instance requiring network switch.

Instance ID/Name	Status	Architecture	Version	Configuration	Network	AZ	Private Network Address	Billing Mode	Project
Search for "Instance ID:mssql-22shj42", 1 result found Back to list									
mssql-2	Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/30 GB Storage Enhanced SSD Dedicated		Shanghai Zone 4	:1433	Monthly Subscription 2024-04-16 20:06:49	Default
						Primary Shanghai Zone 2			
						Replica			

2. In the **Instance Info** on the **Instance Details** page on the right-hand side, click on **Change Network** under the replica node.

3. In the Change Network pop-up window, after modifying the network and read-only address, click **OK**.

Change Network

1. Switching the network will change the instance IP, and the original IP address will be immediately repossessed

2. Only VPC and subnet in the same region where the instance resides can only be selected

Select Network

1

CIDR: 1 '24, subnet IP/available IP: 253/248
In the current network environment, only CVM in the "I" "I" can access this database [Create VPC](#)
[Create Subnet](#)

Read-Only Address ☒ Auto-Assign IP ☐ Specify IP

OK

Cancel

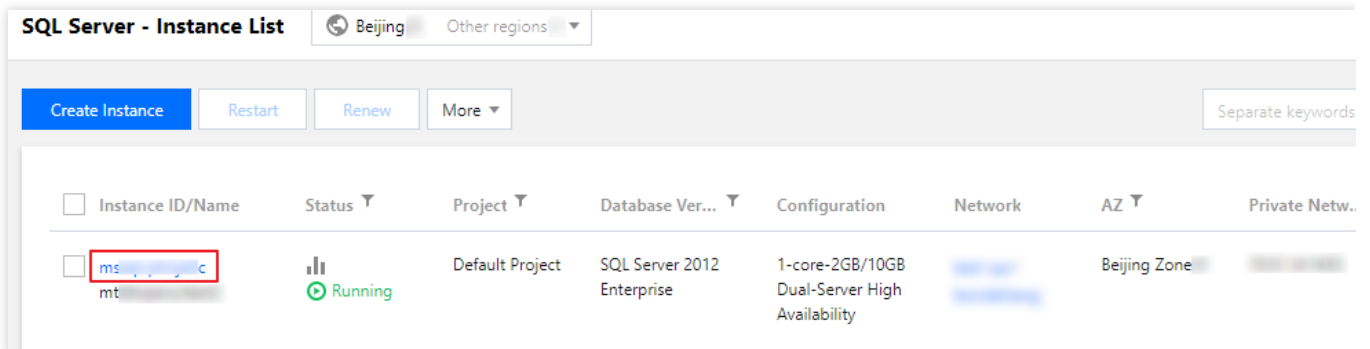
Select Network: After a VPC network is selected, only hosts corresponding to the VPC network can access the database.

Auto-assign IP: A new IP address will be automatically assigned by the system.

Specify IP: You can customize the subnet IP address.

Scenario 2 Operation Steps: Only changing the private network address with network unchanged

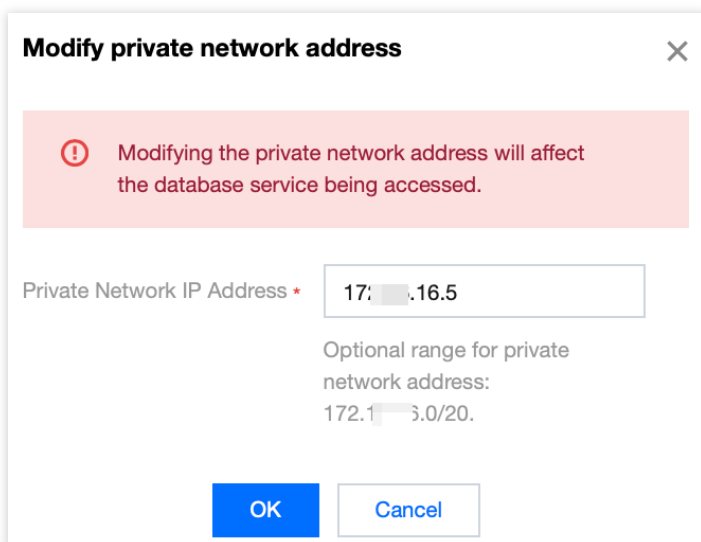
1. Log in to the [SQL Server Console](#), select the region, click the target **Instance ID/Name** or **Manage** in the **Operation** column.



2. In the **Instance Info** on the **Instance Details** page on the right-hand side, click the edit icon following the private network address under the primary node or the current read-only instance.

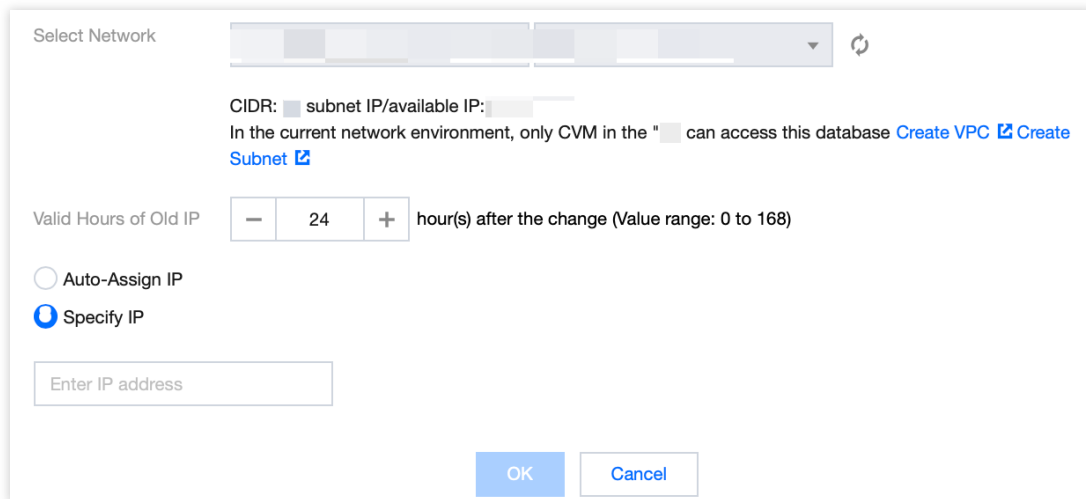


3. In the pop-up window, after modifying the private network address, click **OK**.



Note:

You may also navigate to the **Instance Info** on the **Instance Details** page on the right-hand side, click **Change Network** under the primary node or the current read-only instance. In the pop-up window, you can change the private network address simply by setting the recovery time and specifying an IP, without changing the network itself.



The screenshot shows a 'Select Network' dialog box. At the top, there is a 'Select Network' label and a dropdown menu with a refresh icon. Below this, the 'CIDR' field is shown with a placeholder 'subnet IP/available IP:'. A note states: 'In the current network environment, only CVM in the " " can access this database'. There are two links: 'Create VPC' and 'Create Subnet'. Below the note, there is a section for 'Valid Hours of Old IP' with a minus button, the number '24', and a plus button, followed by the text 'hour(s) after the change (Value range: 0 to 168)'. There are two radio buttons: 'Auto-Assign IP' and 'Specify IP' (which is selected). Below the radio buttons is a text input field labeled 'Enter IP address'. At the bottom right, there are 'OK' and 'Cancel' buttons.

Scenario 3 Operation Steps: Manually releasing reserved addresses

When changing the network, if the recycling time for the old IP address is set to a non-zero hour, the old IP address will be recycled after the set time. You can also manually release the reserved address to release the old IP address in advance.

1. In the **Instance Info** on the **Instance Details** page on the right-hand side, click **Release** following the **Reserved Address**.
2. In the pop-up window, click **OK**, the old IP address will be immediately released. After release, other resources can use the IP address resource in the same subnet.

Enabling/Disabling Public Network Address

Last updated : 2024-08-02 17:27:12

Overview

TencentDB for SQL Server supports both private and public network addresses, with the former enabled by default for you to access your instance over the private network and the latter enabled or disabled as needed.

CLB Architecture Explanation

Currently, when a TencentDB for SQL Server instance is enabled with the public network, it adopts the CLB architecture. The system will automatically create a simple CLB instance of the same region in the [CLB Console](#) to provide public network capability. Please note the resource limitation policy of the CLB architecture (as shown in the table below). If you have higher performance requirements, you can also directly purchase [CLB](#) to achieve this.

Classification	Number of Concurrent Connections	New Connections	Packet Volume	Inbound Bandwidth	Outbound Bandwidth
CLB	2000	200/s	Unlimited	20Mbps	20Mbps

Note:

You can try for free a CLB instance automatically created due to the activation of a public network address. After deactivating the public network address, the corresponding CLB instance in the CLB console will be automatically deleted.

Starting from mid-May 2024, the health probe source IP of CLB will be in the 100.64.0.0/10 segment. If your simple CLB instance shows an abnormal health status after the public network is enabled, you can resolve the issue of health check failure by configuring the security group of your Cloud Database SQL Server instance to **open the 100.64.0.0/10 range**. Please refer to [Configuring Security Groups](#) for the steps.

Inbound rules		Outbound rules			
Add rule		Import rule	Sort by priority	Edit all	Delete
Open all common ports		How to Set		Separate	
<input type="checkbox"/>	Source ?	Protocol:port ?	Policy	Remark	Modific
	Health detection source IP	SQL Server private network IP			
<input type="checkbox"/>	100.64.0.0/10	TCP:1433	Allow	SQL Server(1433) open	2024-0

You need to configure monitoring alerts for the above-mentioned simple CLB instances to monitor public network connections through metrics (such as new public network connections, public network connections, public network outbound bandwidth, and public network inbound bandwidth) after enabling the public network address. Please refer to [Setting Alarm Policies](#) for the operation steps. The policy type is as shown in the image below.

Monitoring Type	Cloud Product Monitoring	HOT APM	HOT RUM	HOT Cloud Probe Monitor
Policy Type	Cloud Load Balancer / Public LB Instance / LB to RS monitor(new) ▼			

Note

After enabling the public network address, you can access your TencentDB for SQL Server instance by using the system-assigned domain name and port. It takes about five minutes for the configuration to take effect.

After the public network access is enabled, it will be controlled by the security group policy. You should configure the database access source in the security group's inbound rules and open the protocol ports (both the private network port (1433 by default) and public network port) as instructed in [Configuring Security Group](#).

Enabling the public network address will expose your database services to the public network, which may lead to database intrusions or attacks. We recommend that you use the private network to connect to the database in the production environment, as public network access may become unavailable due to uncontrollable factors, such as DDoS attacks and large traffic surges.

A public network address makes it less secure to access an instance, and service availability cannot be guaranteed by SLA. Therefore, we recommend that you access your instance at the public network address only when developing, testing, or managing databases. To make transfer faster and ensure a higher security level, use the private network address for database connection. Do not use the public network to sustain the business load, and if you need this, we recommend that you follow the instructions described in [Enabling Public Network Access Through CLB](#).

Currently, enabling the public network address and the resulting traffic are free of charge, but the stability of the public network bandwidth and traffic cannot be guaranteed.

The instance service downtime caused by public network errors won't be counted into the "Single Instance Service Downtime" in TencentDB for SQL Server Service Level Agreement (SLA).

Prerequisites

The instance uses a VPC.

The instance resides in the following regions: Guangzhou, Shanghai, Beijing, Chengdu, Chongqing, Nanjing, Hong Kong (China), Singapore, Seoul, Tokyo, Silicon Valley, or Frankfurt.

Note:

If you can't enable public network access for an instance in the above regions, [submit a ticket](#) for assistance.

Private/Public network address description

Address Type	Description
Private network address	<p>A private network address is an IP address that cannot be accessed by an external device on the internet. It is the implementation form of the Tencent Cloud private network service.</p> <p>A private network address is provided by the system by default and cannot be disabled. You can switch the network type though.</p> <p>If your CVM and TencentDB for SQL Server instances are in the same VPC in the same region under the same Tencent Cloud root account, they can be interconnected over the private network, and there is no need to enable the public network address.</p> <p>It is highly secure.</p>
Public network address	<p>A public network address is a non-reserved address on the internet.</p> <p>A public network address needs to be manually enabled and can be disabled when no longer needed.</p> <p>As a public network address will expose your instance to security risks, it should be used with caution.</p> <p>A device not in Tencent Cloud can access a TencentDB for SQL Server instance at its public network address.</p>

Directions

The procedures for enabling or disabling public network addresses for the primary instance and read-only groups are slightly differently. For the primary instance, configurations must be made within the instance details page. As for read-only groups, configurations must be made within the read-only group of the corresponding primary instance. The following sections will elaborate on the steps respectively.

Note:

The independent enabling or disabling of public network addresses is not supported by read-only instances. It is only supported by the read-only group which the read-only instance belongs to. Moreover, it can only be configured within the read-only group which the read-only instance belongs to, and cannot be set from the details page of the read-only instance.

Enabling/Disabling Public Network Addresses for Primary Instances

Enabling/Disabling Public Network Addresses for Read-Only Groups

Enabling Public IP Addresses for Primary Instances

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region and click the ID or **Manage** in the **Operation** column of the target instance in the instance list.
3. On the **Instance Details** page, click **Enable** in **Basic Info** > **Public Address**.

Basic Info

Instance ID: mssql-
Instance Name: fe0642e-808-61c2a250390d
Instance Remarks: --
Status: Running
Version: SQL Server 2008 R2 Enterprise
Architecture: Two-Node (Local Disk)
Memory: 2GB
Total Capacity: 10GB
Region: South China(Guangzhou)
AZ: Guangzhou Zone 3
Network:
Private Network Address ①: IP: 172.1433
Public Network Address ①: **Enable**

Instance Architecture Diagram

Two-Node instance

Guangzhou/Guangzhou Zone 3

mssql-a fe0642ad-0-61c2a2... (Running) Current Instance

Master-Slave Sync

Replica Guangzhou Zone 3

Backup Center (Guangzhou)

Add Read-Only Instance

Basic Info

Instance ID: mssql-
Instance Name: fe0642e-808-61c2a250390d
Instance Remarks: --
Status: Running
Version: SQL Server 2008 R2 Enterprise
Architecture: Two-Node (Local Disk)
Memory: 2GB
Total Capacity: 10GB
Region: South China(Guangzhou)
AZ: Guangzhou Zone 3
Network:
Private Network Address ①: IP: 172.1433
Public Network Address ①: **Enable**

Instance Architecture Diagram

Two-Node instance

Guangzhou/Guangzhou Zone 3

mssql-a fe0642ad-0-61c2a2... (Running) Current Instance

Master-Slave Sync

Replica Guangzhou Zone 3

Backup Center (Guangzhou)

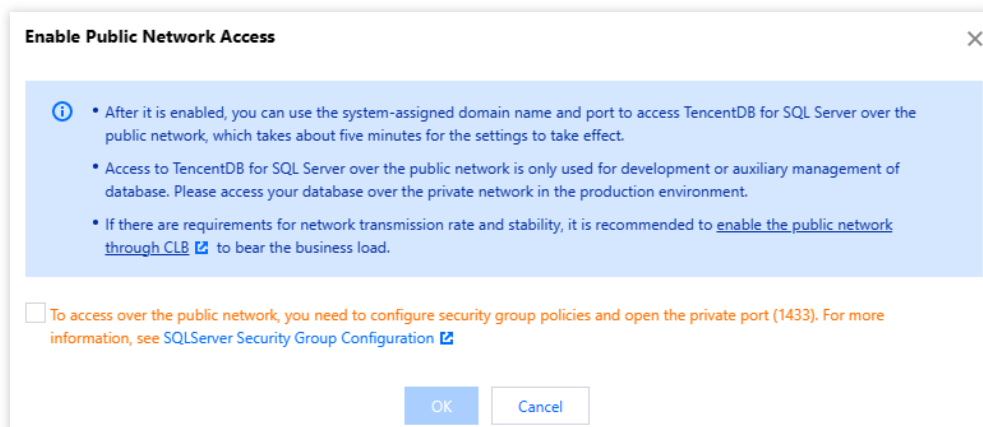
Add Read-Only Instance

4. In the **Enabling public network** window, read the note, indicate your consent, and click **OK** (before the public network address is enabled, a note will be displayed depending on whether a security group is configured).

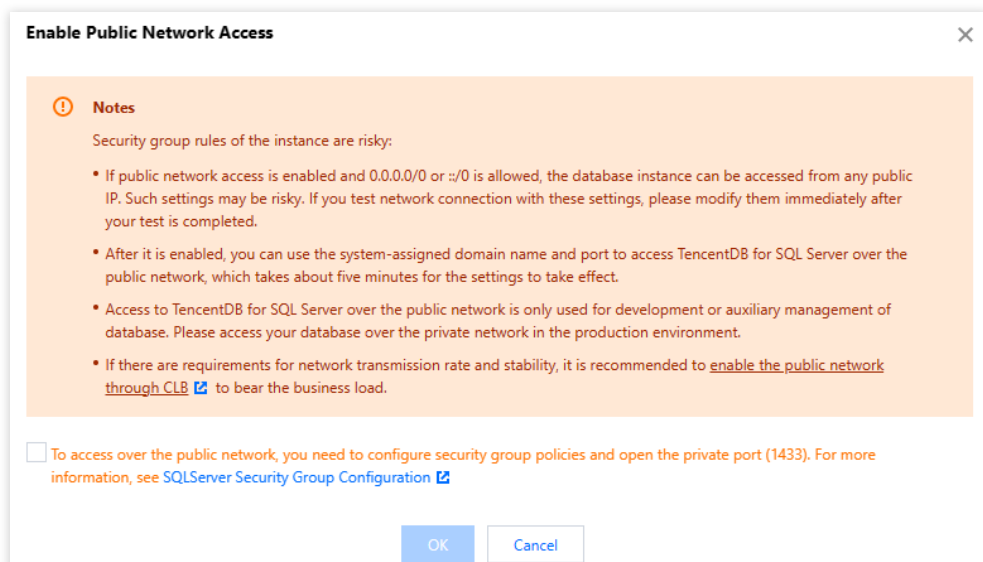
Note:

After the public network address is enabled, it can be viewed in **Basic Info**. The public network access can be toggled off. When it is enabled again, the public network address corresponding to the domain name remains the same.

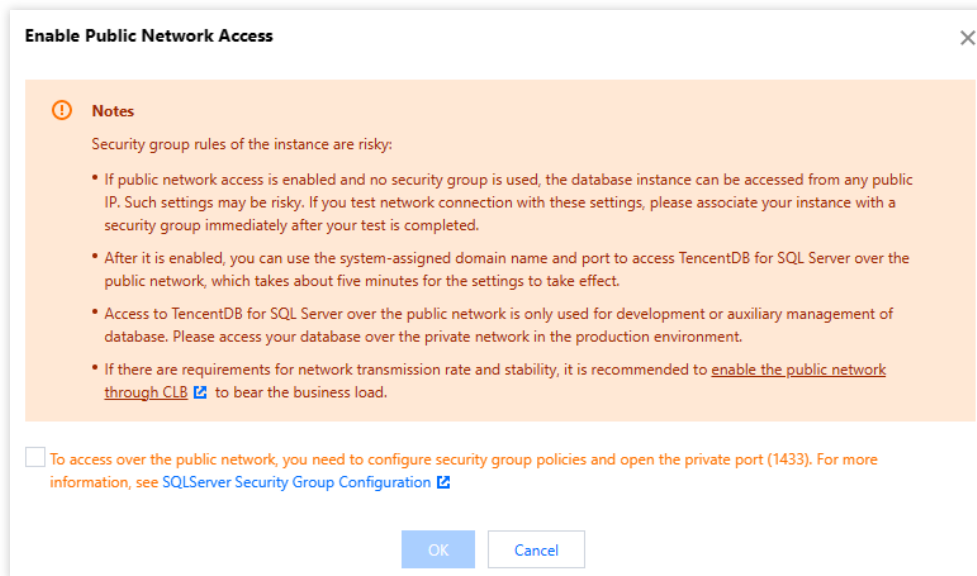
If your instance is bound to a security group, and no high-risk policy is involved, the public network address can be enabled, and a note will be displayed as follows:



If your instance is bound to a security group, but there is a high-risk inbound rule such as `0.0.0.0/0` or `:::/0`, a note will be displayed as follows:



If your instance is not bound to a security group, enabling public network access will lead to a high risk, and a note will be displayed as follows:



5. After the instance status becomes **Running**, you can view the public network address on the instance details page.

Disabling the public network address


1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region and click the ID or **Manage** in the **Operation** column of the target instance in the instance list.
3. On the **Instance Details** tab, click **Disable** in **Basic Info > Public Network Address**.
4. In the **Disabling public network** pop-up window, click **OK**.

Note:

After it is disabled, you can no longer use the domain name and port to access TencentDB for SQL Server over the public network. To minimize potential losses, make sure that no public address is used in your system before disabling it.

Enabling Public Network Addresses for Read-Only Groups

1. Log in to the [SQL Server Console](#).
2. Select the region, and in the instance list, locate the primary instance for which you want to enable the public network address of the read-only group. Click **Instance ID** or **Manage** option in the **Operation** column.
3. Go to the **Read-Only Instance** page from the **Instance Details** page, then click **Enable** following the **Public Network Address** under the **RO group**.



4. In the window for enabling public network settings, read and check the prompt, then click **OK**.

Note:

When the public network address is enabled, it can be viewed in the RO group or the basic information of the corresponding read-only instance. The public network connection can be enabled via the toggle switch. And when you re-enable the public network, the public network address corresponding to the domain remains unchanged.

Disabling Public Network Addresses for Read-Only Groups

1. Log in to the [SQL Server Console](#).
2. Select the region, and in the instance list, locate the primary instance for which you want to disable the public network address of the read-only group. Click **Instance ID** or **Manage** option in the **Operation** column.
3. Go to the **Read-Only Instance** page from the **Instance Details** page, then click **Close** following the **Public Network Address** under the **RO group**.
4. In the pop-up window for disabling public network access, click **OK**.

Note:

After the public network access is disabled, you cannot access the read-only group corresponding to the SQL Server primary instance via the public domain names and ports. Please ensure that your application system does not use public access addresses to avoid losses.

Enabling Public Network Access Through CLB

Last updated : 2024-05-23 21:35:42

TencentDB for SQL Server supports both private and public network addresses, with the former enabled by default for you to access your instance over the private network and the latter enabled or disabled as needed. To access your database instance from a Linux or Windows CVM instance over the public network, you can enable the public network address. You can also enable public network access through CLB, but you must configure security group rules in this case.

This document describes how to bind a TencentDB for SQL Server instance to CLB to enable public network access, connect to the instance through SQL Server Management Studio (SSMS), and run a simple query.

Prerequisites

You have applied for using the backend service feature.

1. Go to the [Cross-Region CLB Binding 2.0 Application page](#).
2. Fill out and submit the application.
3. Upon completion of the beta application submission, [submit a ticket to CLB](#) to apply for access to the backend service features.

Note:

Enabling the public network services through CLB is exclusively applicable when the CLB instances and the SQL Server instances belong to the same VPC network. It is currently not supported for instances across different VPC networks.

Step 1. Purchase a CLB instance

Note:

If you already have a CLB instance in the same region as TencentDB for SQL Server, skip this step.

Go to the [CLB purchase page](#), select the configuration, and click **Buy Now**.

Note:

Region: You need to select the region where the TencentDB for SQL Server instance is.

Step 2. Configure the CLB instance

The following describes how to configure the CLB instance in the same VPC as the database instance and in a different VPC respectively.

Scenario: Deploying the CLB instance in the same VPC as the TencentDB for SQL Server instance

1. Enable cross-VPC access so that the CLB instance can be bound to another private IP.
 - a. Log in to the [CLB console](#), select the region, and click the target instance ID in the instance list to enter the instance management page.
 - b. On the **Basic Info** page, click **Configure** in the **Real Server** section.
 - c. In the pop-up window, click **Submit**.
2. Configure a public network listener port.
 - a. Log in to the [CLB console](#), select the region, and click the target instance ID in the instance list to enter the instance management page.
 - b. On the instance management page, select the **Listener Management** tab and click **Create** below **TCP/UDP/TCP SSL Listener**.
 - c. In the pop-up window, complete the settings and click **Submit**.

Step 3. Bind a TencentDB for SQL Server instance

1. After creating the listener, click it in **Listener Management** and click **Bind** on the right.
2. In the pop-up window, select **Other Private IPs** as the object type, enter the IP address and port of the TencentDB for SQL Server instance, and click **OK**.

Note:

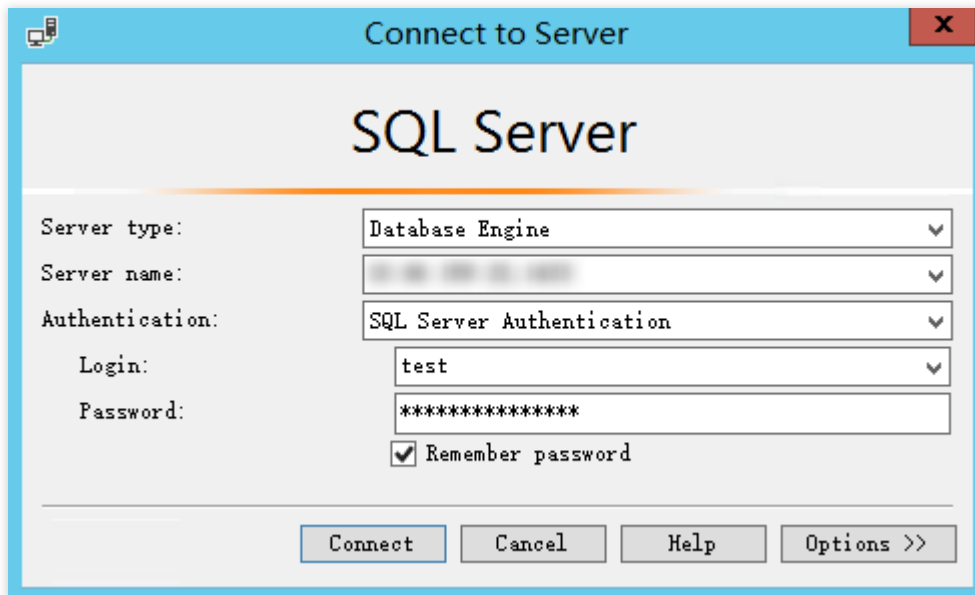
The login account must be a standard account (bill-by-IP). If binding fails, [submit a ticket](#) for assistance.

Step 4. Configure the TencentDB for SQL Server security group

1. Log in to the [TencentDB for SQL Server console](#), select a region, and click the ID of the target instance in the instance list or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Security Group** tab, click **Configure Security Group**, configure the security group rule to open all ports, and confirm that the security group allows access from public IPs. For more information on configuration, see [Configuring Security Group](#).

Step 5. Connect to the TencentDB for SQL Server instance over the public network

1. Download and install [SSMS](#) locally. For more information on SSMS, see [What is SQL Server Management Studio \(SSMS\)?](#)
2. Start SSMS locally. On the **Connect to Server** page, enter the relevant information to connect to TencentDB. Click **Connect** and wait a few minutes before SSMS connects to your database instance.



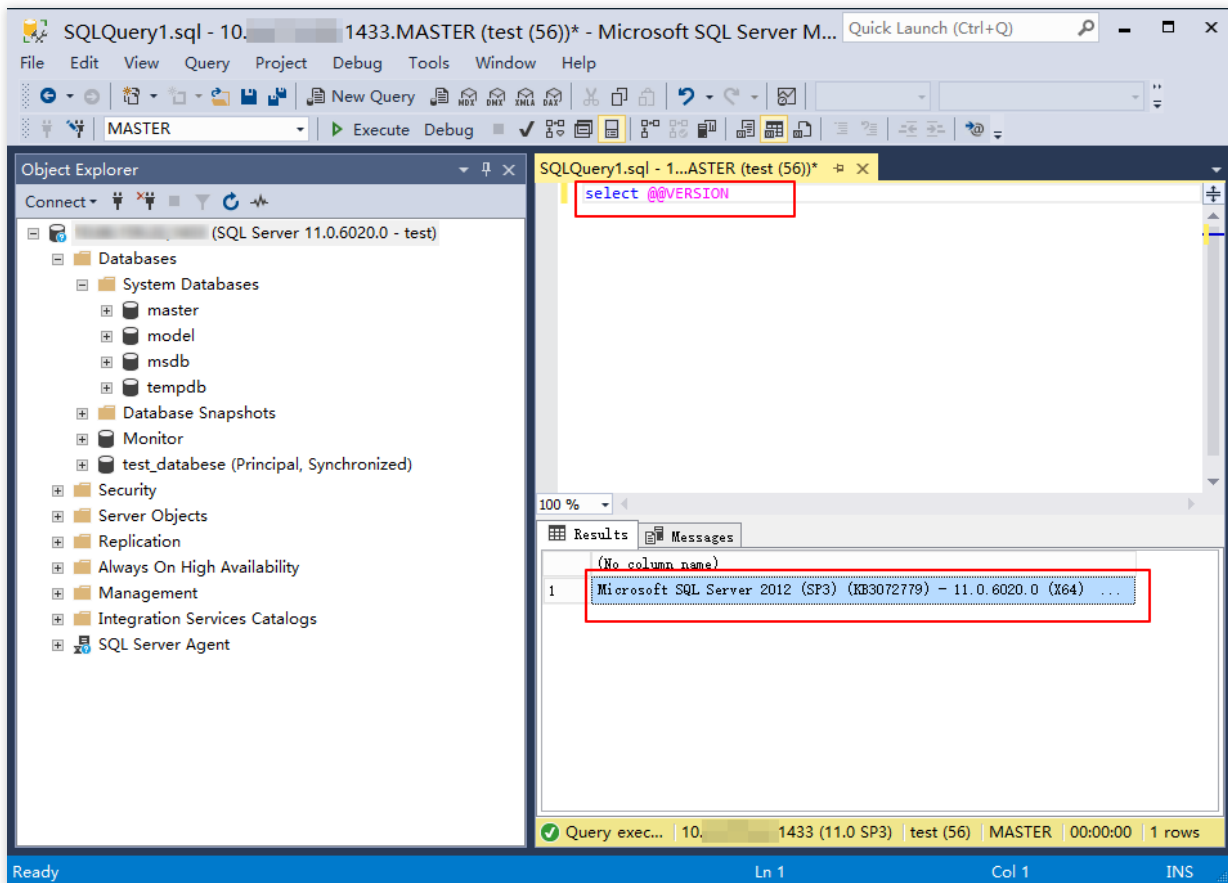
Server type: Select **Database Engine**.

Server name: Enter the local IP address and port number of the CLB instance and separate them by a comma, such as `10.0.0.1,4000`.


Authentication: Select **SQL Server Authentication**.

Login and **Password:** Enter the account name and password you configured when creating the instance account on the **Account Management** tab.

3. Once connected to the database, you can view the standard built-in system databases (master, model, msdb and tempdb) of SQL Server.

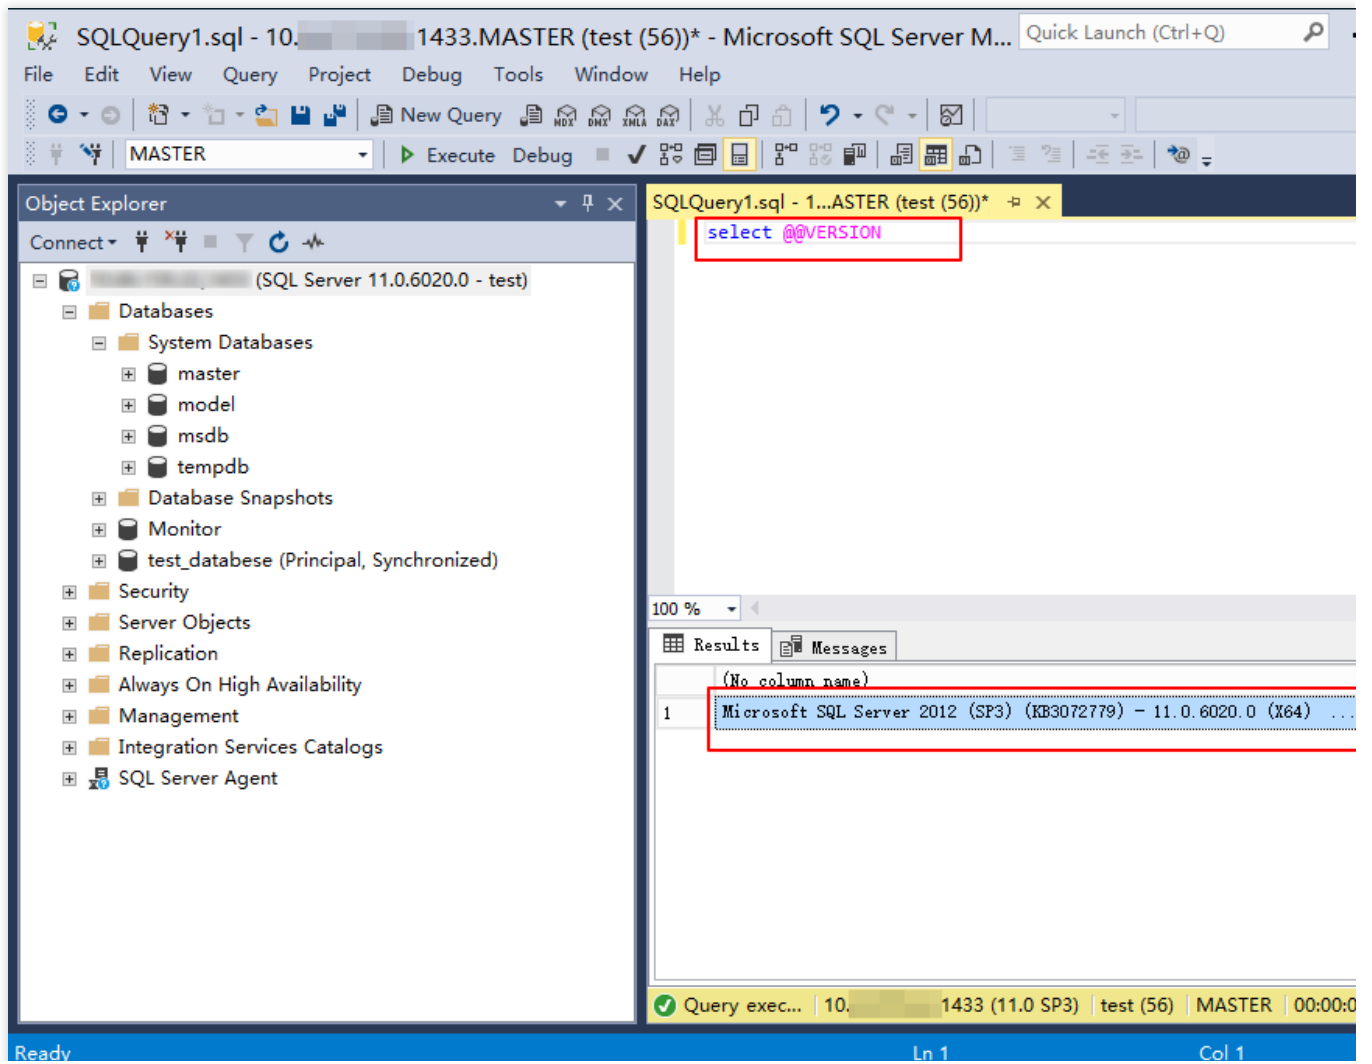


4. Now you can start creating your own databases and running queries for them. Select **File > New > Query with Current Connection** and type the following SQL query:



```
select @@VERSION
```

Run the query. SSMS returns the SQL Server version of TencentDB instance.



CAM

Overview

Last updated : 2024-01-18 17:20:33

Known Issues

If you use multiple Tencent Cloud services such as CVM, VPC, and TencentDB that are managed by different users sharing your Tencent Cloud account key, you may face the following problems:

Your key is shared by multiple users, leading to high risk of compromise.

You cannot limit the access permissions of other users, which poses a security risk due to potential faulty operations.

Solution

You can allow different users to manage different services through sub-accounts so as to avoid the above problems. By default, a sub-account doesn't have permission to use a Tencent Cloud service or related resources. Therefore, you need to create a policy to grant the required permission to the sub-account.

[Cloud Access Management \(CAM\)](#) is a web-based Tencent Cloud service that helps you securely manage and control access permissions to your Tencent Cloud resources. Using CAM, you can create, manage, and terminate users (groups), and control the Tencent Cloud resources that can be used by the specified user through identity and policy management.

When using CAM, you can associate a policy with a user or user group to allow or forbid them to use specified resources to complete specified tasks. For more information on CAM policies, please see [Policy Syntax](#).

If you do not need to manage the access permissions to TencentDB resources for sub-accounts, you can skip this chapter. This will not affect your understanding and usage of other parts in the documentation.

Getting started

A CAM policy must authorize or deny the use of one or more TencentDB operations. At the same time, it must specify the resources that can be used for the operations (which can be all resources or partial resources for certain operations). A policy can also include the conditions set for the manipulated resources.

You are recommended to manage TencentDB resources and authorize TencentDB operations through CAM policies. Although the experience stays the same for existing users who are granted permission by project, it is not recommended to continue managing resources and authorizing operations in a project-based manner. Effectiveness conditions cannot be set in TencentDB for the time being.


Task	Link
Basic policy structure	Policy Syntax
Operation definition in a policy	TencentDB Operations
Resource definition in a policy	TencentDB Resource Path
Resource-level permission supported by TencentDB	Resource-level Permission Supported by TencentDB

Authorization Policy Syntax

Last updated : 2024-01-18 17:20:33

Policy Syntax

CAM policy:



```
{
```

```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "effect",
      "action": ["action"],
      "resource": ["resource"],
      "condition": {"key": {"value": {}}}
    }
  ]
}
```

version is required. Currently, only "2.0" is allowed.

statement describes the details of one or more permissions. This element contains a permission or permission set of other elements such as effect, action, resource, and condition. One policy has only one statement.

effect describes whether the result produced by the statement is "allowed" (allow) or "denied" (deny). This element is required.

action describes the allowed or denied action (operation). An operation can be an API (prefixed with "sqlserver:"). This element is required.

resource describes the details of authorization. A resource is described in a six-segment format. Detailed resource definitions vary by product. This element is required.

condition describes the condition for the policy to take effect. A condition consists of operator, action key, and action value. A condition value may contain information such as time and IP address. Some services allow you to specify additional values in a condition. This element is required.

Operations in TencentDB for SQL Server

In a TencentDB for SQL Server policy statement, you can specify any API operation from any service that supports TencentDB for SQL Server. APIs prefixed with `sqlserver:` should be used for TencentDB for SQL Server, such as `sqlserver:DescribeDBInstances` or `sqlserver:CreateAccount`.

To specify multiple operations in a single statement, separate them with commas, as shown below:




```
"action":["sqlserver:action1","sqlserver:action2"]
```

You can also specify multiple operations using a wildcard. For example, you can specify all operations beginning with "Describe" in name, as shown below:



```
"action":["sqlserver:Describe*"]
```

If you want to specify all operations in TencentDB for SQL Server, use a wildcard as shown below:



```
"action": ["sqlserver:*"]
```

TencentDB for SQL Server Resources

Each CAM policy statement has its own resources.

Resources are generally in the following format:



```
qcs:project_id:service_type:region:account:resource
```

project_id describes the project information, which is only used to enable compatibility with legacy CAM logic and can be left empty.

service_type describes the product abbreviation such as sqlserver.

region describes the region information, such as ap-guangzhou.

account is the root account of the resource owner, such as uin/653339763.

resource describes detailed resource information of each product, such as instance/instance_id1 or instance/*.

For example, you can specify a resource for a specific instance (mssql-m8oh024t) in a statement as shown below:



```
"resource": [ "qcs::sqlserver:ap-guangzhou:uin/653339763:instance/mssql-m8oh024t"]
```

You can also use the wildcard "*" to specify it for all instances that belong to a specific account as shown below:




```
"resource": [ "qcs::sqlserver:ap-guangzhou:uin/653339763:instance/*"]
```

If you want to specify all resources or a specific API operation does not support resource-level permission control, you can use the wildcard "*" in the "resource" element as shown below:



```
"resource": ["*"]
```

To specify multiple resources in a single command, separate them with commas. Below is an example where two resources are specified:



```
"resource":["resource1","resource2"]
```

The table below describes the resources that can be used by TencentDB for SQL Server and the corresponding resource description methods, where words prefixed with \$ are placeholders, `project` refers to a project ID, `region` refers to a region, and `account` refers to an account ID.

Resource	Resource Description Method in Authorization Policy
Instance	<code>``qcs::sqlserver:\$region:\$account:instance/\$instanceId``</code>
VPC	<code>``qcs::vpc:\$region:\$account:vpc/\$vpcId``</code>

Security group	<code>`qcs::cvm:\$region:\$account:sg/\$sgId`</code>
----------------	--

Authorizable Resource Types

Last updated : 2024-01-18 17:20:33

Resource-level permission can be used to specify which resources a user can manipulate. TencentDB for SQL Server supports certain resource-level permissions. This means that for TencentDB for SQL Server operations that support resource-level permission, you can control the time when a user is allowed to perform operations or to use specified resources. The following table describes the types of resources that can be authorized in CAM.

Resource Type	Resource Description Method in Authorization Policy
TencentDB instance-related resource	<code>qcs::sqlserver:\$region:\$account:instance/*</code> <code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code>

TencentDB for SQL Server supports resource-level authorization. You can allow a sub-account to have API permissions for specific resources. The table below lists the TencentDB API operations which currently support resource-level permission control as well as the resources and condition keys supported by each operation. When specifying a resource path, you can use the "*" wildcard in the path.

Any TencentDB API operation not listed in the table does not support resource-level permission. For such an operation, you can still authorize a user to perform it, but you must specify `*` as the resource element in the policy statement.

API Name	Description	Six-segment Example of Resource
CreateAccount	Creates an account	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
CreateBackup	Creates a backup	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
CreateDB	Creates a database	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DeleteAccount	Deletes an account	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DeleteDB	Drops a database	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeAccounts	Queries the	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code>

	account list	<code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeBackups	Queries the backup list	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeDatabaseNames	Queries database names	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeDBInstances	Queries the instance list	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeDBs	Queries the database list	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeInstanceTasks	Queries instance tasks	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeRollbackTime	Queries the time range available for rollback	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
DescribeSlowlogs	Queries the slow log list	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
InquiryPriceRenewDBInstance	Queries the renewal price of an instance	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
InquiryPriceUpgradeDBInstance	Queries the upgrade price of an instance	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyAccountPrivilege	Modifies account permissions	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyAccountRemark	Modifies account remarks	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>``qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyBackupStrategy	Modifies the	<code>``qcs::sqlserver:\$region:\$account:instance/\$`</code>

	time for cold backup	<code>`qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyDatabasePrivilege	Modifies database permissions	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyDBInstanceName	Renames an instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyDBInstanceProject	Modifies an instance project	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyDBName	Renames a database	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
ModifyDBRemark	Modifies database remarks	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
RenewDBInstance	Renews an instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
ResetAccountPassword	Resets an account password	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
RestartDBInstance	Restarts an instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
RestoreInstance	Restores a cold backup instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
RollbackInstance	Rolls back an instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
TerminateDBInstance	Terminates an instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>
UpgradeDBInstance	Upgrades an instance	<code>`qcs::sqlserver:\$region:\$account:instance/\$`</code> <code>`qcs::sqlserver:\$region:\$account:instance/`</code>

Account Management

Account Types and Permissions

Last updated : 2024-01-18 17:20:33

After you create a TencentDB for SQL Server instance, you can create different database accounts to allocate and manage databases based on your business needs.

You can create different types of accounts with different permissions for both TencentDB for SQL Server two-node (formerly High Availability/Cluster Edition) and single-node (formerly Basic Edition) instances. This document describes the supported types of accounts and their permissions.

Note:

TencentDB for SQL Server launched the new database account and permission logic on February 9, 2023. For the mappings between old and new account types and permissions, see [Account Type and Permission Changes](#).

Account types and permissions for two-node (formerly High Availability/Cluster Edition) instances

Instance Architecture	Account Type	Database Permission	Role Description
Two-node (formerly High Availability/Cluster Edition)	Privileged account	Instance admin account, which has the owner permissions of all databases by default.	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_owner
	Standard account	Owner	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_owner
		Read/Write	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_reader db_writer

		Read-only	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_reader
	Designated account	A designated account can only view and own the specified database. A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_owner

Account types and permissions for single-node (formerly Basic Edition) instances

Instance Architecture	Account Type	Database Permission	Role Description
Single-node (formerly Basic Edition)	Admin account	Instance admin account, which has the highest-level sysadmin permission and the owner permissions of all databases. After the admin account is enabled, the product SLA will no longer be guaranteed.	Server-level roles: sysadminDatabase-level roles: db_owner
	Privileged account	It has the owner permissions of all databases by default.	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_owner
	Standard account	Owner	Server-level roles: securityadmin processadmin dbcreatorDatabase-level roles: db_owner

		Read/Write	Server-level roles: securityadmin processadmin dbcreator Database-level roles: db_reader db_writer
		Read-only	Server-level roles: securityadmin processadmin dbcreator Database-level roles: db_reader
	Designated account	A designated account can only view and own the specified database. A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.	Server-level roles: securityadmin processadmin dbcreator Database-level roles: db_owner

Account List

Last updated : 2024-01-18 17:20:33

This document describes the fields in the account list.

Database account


After you create a TencentDB for SQL Server instance, you can create different database accounts to allocate and manage databases based on your business needs.

Viewing the account list

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Manage** in the **Operation** column of the target instance to enter the **Instance Details** page.
3. Click **Account Management** to enter the corresponding tab.

Instance Details	System Monitoring	Backup and Restoration	Security Group	Account Management	Database Management	Bac
<div>Create AccountBatch Management</div>						
<input type="checkbox"/>	Account Name	Account Type	Status	Account Creation T... ⚙	Account Update Ti... ⚙	Password Update T... ⚙ Da
<input type="checkbox"/>	tor	Standard Account	Normal	2023-02-07 17:20:05	2023-02-07 17:20:46	2023-02-07 17:20:46 --

Tools

Tool	Description
Search box	Click  to filter accounts. You can perform a fuzzy search for accounts by account name.
Refresh	You can click



to refresh the list.

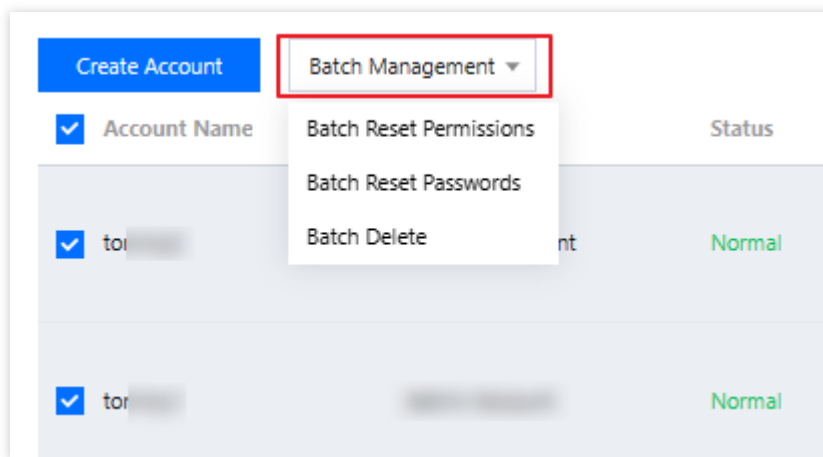
Fields in the account list

Parameter	Description
Account Name	The names of the accounts created for the instance.
Account Type	<p>The types of accounts created for the instance. The supported account types vary by instance architecture as follows:</p> <p>Two-node</p> <p>Privileged Account: A type of admin account, which has owner privileges for all databases by default.</p> <p>Standard Account: A type of non-admin account, which supports specifying readwrite, read-only, and owner privileges for a specific database.</p> <p>Designated Account: A designated account can only view and own the specified database. A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.</p> <p>Single-node</p> <p>Admin Account: A type of admin account, which has the highest-level <code>sysadmin</code> permission and the owner permissions of all databases. After it is enabled, the product SLA will no longer be guaranteed.</p> <p>Privileged Account: A type of admin account, which has owner privileges for all databases by default.</p> <p>Standard Account: A type of non-admin account, which supports specifying readwrite, read-only, and owner privileges for a specific database.</p> <p>Designated Account: A designated account can only view and own the specified database. A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.</p>
Status	The status of the account.
Account Creation Time	The time when the account was created, which can be sorted in ascending or descending order.
Account Update Time	The time when the account was last updated, which can be sorted in ascending or descending order.
Password Update Time	The time when the account password was last reset, which can be sorted in ascending or descending order.

Database	The databases that the account is authorized to access.
Remarks	The remarks of the account.
Operation	Set Permissions: Set the read/write or read-only permission of the specified database for the account. Reset Password: Reset the password of the account. Delete Account: Delete the account.

Note:

For each operation in the account list, you can select multiple accounts and click **Batch Management** above the list for batch management.



Relevant operations

[Creating Account](#)[Modifying Account Permissions](#)[Resetting Password](#)[Deleting Account](#)

Creating Account

Last updated : 2024-01-18 17:20:33

Overview

TencentDB for SQL Server supports creating and deleting accounts and modifying account permissions on the **Manage Account** page in the console. Such operations cannot be performed in Microsoft SQL Server Management.

Note:

The created account name and password will be used when connecting to TencentDB for SQL Server. Please store them properly.

Directions

1. Log in to the [TencentDB for SQL Server console](#) and click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select **Manage Account** > **Create Account** and enter relevant information in the pop-up window. After confirming that everything is correct, click **OK**.

Account Name: it is required, can contain 1–50 letters, digits, or underscores, and must start with a letter.

Admin: it is optional. Each TencentDB for SQL Server instance can have one admin account, which has read/write permissions to all databases by default. A newly added database will be automatically authorized for the admin account with no manual authorization required.

Note:

There are three types of account permissions in TencentDB for SQL Server:

Admin: by default, it has read/write permissions to all databases. Only one admin account can be set for one instance.

Read/Write: it has read/write permissions to authorized databases and can perform database changes.

Read-only: it has read-only permission to only authorized databases and cannot perform changes.

Database: it is optional. You can set the permissions (read/Write or read-only) that the account has to the database when creating an account. You can also grant permissions when modifying permissions or creating databases.

Password: it is required and should be a combination of 8–32 characters comprised of at least two of the following types: letters, digits, and special symbols (`_+-&=!@#$%^*()[]`).

Account Remarks: it is optional and can contain up to 256 characters.

Create Account

Account Name

Account name cannot be empty

Admin

Database

Unauthorized Database

Authorized Dat...Permissions

New Password

Enter the password

At least 8 bits for password

Confirm Password

Please enter confirm password

Remarks

Enter remarks

Up to 256 chars for remarks

OK

Cancel

3. After the account is created, you can perform operations such as **Remove Admin**, **Modify Permissions**, **Reset Password**, and **Delete Account** in the account list.

Create AccountBatch Management

<input type="checkbox"/>	Account Name	Status	Creation Time	Update Time	Password Update Time	Database	Remarks
<input checked="" type="checkbox"/>	weifg	Creating	2021-03-25 19:25:29	2021-03-25 19:25:29	--	--	--

Deleting Account

Last updated : 2024-08-14 10:44:53

Overview

To disable a created database account, you can delete it in the TencentDB for SQL Server console.

Note:

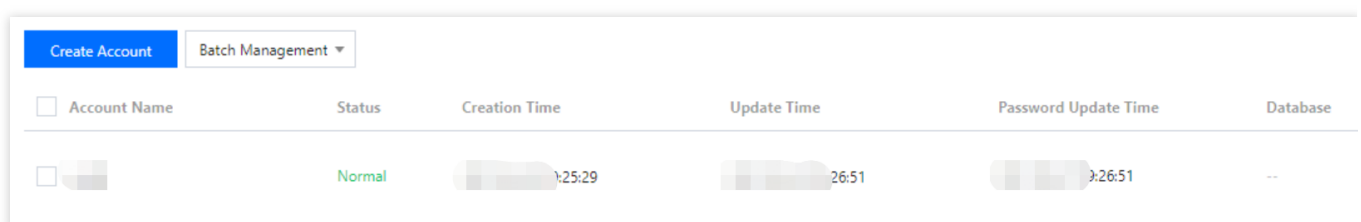
In order to avoid accidental deletion from interrupting normal use by your business, you need to make sure that the database account to be deleted is no longer used by any applications.

Instructions for Sensitive Operations with MFA Integrated

To enhance the security of the cloud account, TencentDB for SQL Server supports MFA (Multi-Factor Authentication), which brings an extra layer of protection in addition to the username and password. After the MFA device verification is enabled, when you perform operations such as destroying instances/changing network/modifying IP addresses/resetting passwords/deleting accounts/resetting instances/performing primary-replica switch, a second identity verification based on the MFA dynamic code will be conducted. Such operations can be executed only after successful verification. For an introduction to MFA and how to enable the operation protection, see [MFA Devices](#).

Directions

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Manage Account** tab, select the target account, and click **Delete Account** in the **Operation** column.



<input type="checkbox"/>	Account Name	Status	Creation Time	Update Time	Password Update Time	Database
<input type="checkbox"/>		Normal	1:25:29	26:51	3:26:51	--

3. In the pop-up window, confirm the information and click **Delete**.

Modifying Account Permissions

Last updated : 2024-01-18 17:20:33

Overview

TencentDB for SQL Server supports modifying account permissions as needed. However, permissions of the admin account cannot be changed to specified permissions.

Directions

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Account Management** tab, select the desired account, and click **Set Permissions** in the **Operation** column.

Create Account		Batch Management ▾				
<input type="checkbox"/>	Account Name	Status	Creation Time	Update Time	Password Update Time	Database
<input type="checkbox"/>	werfg	Normal	2021-03-25 19:25:29	2021-03-25 19:26:51	2021-03-25 19:26:51	--

3. In the pop-up window, select or deselect the authorized databases or permissions and click **OK** to complete the modification.

Set Permissions

Database Account

werfg

Admin

Database List

Unauthorized Database

Authorized Dat...

Permissions

↔

OK

Cancel

Creating Admin Account

Last updated : 2024-01-18 17:20:33

This document describes how to create an admin account for a TencentDB for SQL Server instance in the console.

Admin account description

The System Admin (SA) role is the most privileged role in SQL Server. It can perform any operations in SQL Server without passing any security checks.

TencentDB for SQL Server single-node (formerly Basic Edition) instances allow you to create an admin account with the SA privilege, so you can use the account to quickly adapt offline software and move it to the cloud.

Warning:

Once an admin account is created for an instance, the `sysadmin` privilege is enabled, and the product SLA will no longer be guaranteed for the instance. Database instances that do not have admin accounts with the SA privilege are not affected.

You can create only one SA account for each instance.

Directions

1. Log in to the [TencentDB for SQL Server console](#). Click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the instance management page, select **Account Management** > **Create Account** and enter relevant information in the pop-up window. After confirming that everything is correct, click **OK**.

Create Account

i We will release a new version of database account types and permissions for TencentDB for SQL Server on February 10, 2023. For details about the mappings between the old and new versions, see [Account Type and Permission Modification](#).

Account Name

chaojiqianxian

The account name must contain 1-50 letters, digits, or underscores, and start with a letter.

Account Type

☒ Admin Account ☐ Privileged Account ☐ Standard Account ☐ Designated Account

As the admin account has the highest-level management permissions, the instance SLA will not be guaranteed after this account is enabled.

☒ I have read the risks of [Creating System Admin Account](#) and agreed to [Service Level Agreement](#)

Database

New Password

The password must contain 8–32 characters in at least two of the following three types: letters, digits, and special characters (including underscore, hyphen, equals, exclamation mark, at sign, hash, dollar sign, percent, caret, asterisk, and open and close square brackets).

Confirm Password

Remarks

Enter remarks

Up to 256 characters for remarks

OK

Cancel

Parameter	Description
Account Name	It can contain 1–50 letters, digits, or underscores and must start with a letter. The account name cannot contain the following keywords or symbols: <code>sysadmin</code> , <code>sp_addsrvrolemember</code> , <code>master</code> , <code>mssql</code> , <code>##</code> , <code>[]</code> , <code>,</code> , <code>.</code> , <code>;</code> , <code>()</code> , <code>--</code> .
Account Type	Select Admin Account . Warning: As the admin account has the highest-level management permissions, the instance SLA will not be guaranteed after this account is enabled.

Database	The admin account has the owner permissions of all databases by default.
New Password	The password must contain 8-32 characters in at least two of the following types: letters, digits, and symbols _+&=!@\$%^*()[].
Confirm Password	Enter the password again.
Remarks	Optional. It can contain up to 256 characters.

Creating Designated Account

Last updated : 2024-01-18 17:20:33

This document describes how to create a designated account for a TencentDB for SQL Server instance in the console.

Designated account description

You can create one or multiple designated accounts for a TencentDB for SQL Server single-node (formerly basic edition) or two-node (formerly high availability/cluster edition) instance. Note the following for designated accounts: A designated account can only view and own the specified database.

For example, account B is created in instance A and designated to have the owner permission of the database db1. Then, when logging in to instance A with designated account B, you can only view the database db1, but you can perform all operations on it.

A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.

This means that a designated account can have the owner permissions of multiple databases, but a database with owner permission granted to a designated account cannot be authorized to other designated accounts.

Directions

1. Log in to the [TencentDB for SQL Server console](#). Click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the instance management page, select **Account Management** > **Create Account** and enter relevant information in the pop-up window. After confirming that everything is correct, click **OK**.

Parameter	Description
Account Name	It can contain 1–50 letters, digits, or underscores and must start with a letter. The account name cannot contain the following keywords or symbols: <code>sysadmin</code> , <code>sp_addsrvrolemember</code> , <code>master</code> , <code>mssql</code> , <code>##</code> , <code>[]</code> , <code>,</code> , <code>o</code> , <code>.</code> , <code>;</code> , <code>()</code> , <code>--</code> .
Account Type	Designated account description Note: A designated account can only view and own the specified database.
Database	Select target databases in the Unauthorized Database list. The selected databases and

	permissions are displayed under Authorized Database on the right, and the selected databases can be deselected.
New Password	The password must contain 8–32 characters in at least two of the following types: letters, digits, and symbols _+-&=!@#\$%^*()[].
Confirm Password	Enter the password again.
Remarks	Optional. It can contain up to 256 characters.

Resetting Password

Last updated : 2024-08-14 10:49:43

Overview

If you forgot your database account password or need to change it while using TencentDB for SQL Server, you can reset it in the console.

Note:


For TencentDB for SQL Server, the password resetting feature has been connected to [CAM](#); therefore, we recommend you exercise tighter control over the permission to the password resetting API or sensitive resources of TencentDB for SQL Server instances by granting such permission only to appropriate personnel. For data security, we recommend you regularly reset the password at least once every three months.

Instructions for Sensitive Operations with MFA Integrated

To enhance the security of the cloud account, TencentDB for SQL Server supports MFA (Multi-Factor Authentication), which brings an extra layer of protection in addition to the username and password. After the MFA device verification is enabled, when you perform operations such as destroying instances/changing network/modifying IP addresses/resetting passwords/deleting accounts/resetting instances/performing primary-replica switch, a second identity verification based on the MFA dynamic code will be conducted. Such operations can be executed only after successful verification. For an introduction to MFA and how to enable the operation protection, see [MFA Devices](#).

Directions

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Account Management** tab, select the target account, and click **Reset Password** in the **Operation** column.



3. In the pop-up window, enter and confirm the new password and click **OK**.

Note:

The password should be a combination of 8–32 chars comprised of at least two of the following types: letters, digits, and special symbols (`_+-&=!@#$$%^()`).

Setting Account CAM Verification

Last updated : 2024-08-05 16:32:27

This document introduces the instructions and operations for setting account CAM verification through the console.

Note:

If you need to enable account CAM verification, [submit a ticket](#) to apply for the allowlist features.

Background

In scenarios where cloud databases are used, it is often necessary to create separate accounts and passwords for the databases and grant access and operation permissions to corresponding users. This method of account management is complex and prone to security issues such as account and password leaks. Based on this background, TencentDB for SQL Server supports the CAM verification feature for accounts. By connecting sub-accounts of Tencent Cloud platform with database accounts and adding CAM credential authentication, the complexity of account permission management is simplified, therefore enhancing database security and account management efficiency.

Overview

If you have high security requirements, you can use this feature to bind CAM with database accounts for verification. You can obtain the corresponding password when requesting to access the database, thereby enhancing database security. It is recommended that CAM verification be enabled in the following two scenarios.

Using CAM verification as a verification mechanism for temporary, individual access to the database.

Using CAM verification as a verification mechanism only for workloads that can be easily retried.

Notes

Use long connections to access the database whenever possible.

Before enabling CAM verification, ensure that the related CAM permission rules are configured in advance.

After enabling CAM verification, you will not be able to change the password.

After disabling CAM verification, you will not be able to obtain access credentials through CAM. Therefore, you need to enter a new password when disabling CAM verification.

Feature Limits

It is recommended to enable CAM verification for no more than 10 accounts within a single instance.

After CAM verification is enabled, the password reset operation for this account is not supported.

Only an account with a single server address is supported to enable CAM verification.

CAM verification cannot be enabled repeatedly for the same account.

The root account does not support CAM verification.

Prerequisites

The [TencentDB for SQL Server instance has been created](#).


The [ticket has been submitted](#) to apply for this feature.

The instance is running.

Step 1: Configure CAM permission rules

Before using the CAM verification feature with the account, you need to configure the related CAM permission rules.

Policy Content



```
{
  "statement": [
    {
      "action": [
        "cam:BuildDataFlowAuthToken"
      ],
      "effect": "allow",
      "resource": [
        "qcs::cam::uin/<User uin>:resourceUser/<Instance ID>/<Username>",
      ]
    }
  ]
}
```

```
],  
  "version": "2.0"  
}
```

<User uin>: Replace with the actual account ID.

<Instance ID>: Replace with the actual instance ID to be authorized.

<Username>: Replace with the actual username to be authorized.

Operation Instructions

1. Log in to the CAM console with the admin account. On the [Policies](#) page, create a custom policy using the Policy Generator (refer to [Creating Custom Policy](#)).

The screenshot shows the 'Visual Policy Generator' interface in the Tencent Cloud CAM console. It is divided into two steps: '1 Edit Policy' and '2 Associate User/User Group/Role'. The 'Visual Policy Generator' tab is active, showing a policy configuration for 'TencentDB for SQL Server (All actions)'. The 'Effect' is set to 'Allow'. The 'Service' is 'TencentDB for SQL Server (sqlserver)'. The 'Action' is 'All actions (*)'. The 'Resource' is set to 'Specific resources'. A warning message states: 'The selected actions include operation-level APIs. If you select this option, the authorization rules for specific resources won't take effect for such A'. A checkbox 'Do no subdivide an API' is checked. Two resource descriptions are shown: 'qcs::sqlserver::uin/2[redacted]:66:instanceid/mssql-jie-...' and 'qcs::sqlserver::uin/2[redacted]:6:user/jnhe'. Both have 'Add a six-segment resource description to restrict the access.' links. The 'Condition' section has a 'Source IP' checkbox and an 'Add other conditions' link. At the bottom, there is a 'Next' button and a character count: 'Characters: 302 (up to 6,144)'.

1 Edit Policy > 2 Associate User/User Group/Role

Visual Policy Generator JSON

▼ TencentDB for SQL Server(All actions)

Effect * ☒ Allow ☐ Deny

Service * TencentDB for SQL Server (sqlserver)

Action * All actions (*)

Resource * [Collapse](#)

☐ All resources ☒ Specific resources

The selected actions include operation-level APIs. If you select this option, the authorization rules for specific resources won't take effect for such A

☒ Do no subdivide an API ⓘ

instanceid qcs::sqlserver::uin/2[redacted]:66:instanceid/mssql-jie-... Edit Delete ☐ Any resource of this type

[Add a six-segment resource description to restrict the access.](#)

instance Specify a instance six-segment resource description for DescribeBackupByFlowId and 114 other action(s). ⓘ ☐ Any resource c

[Add a six-segment resource description to restrict the access.](#)

qcs::sqlserver::uin/2[redacted]:6:user/jnhe Edit Delete ☐ Any resource of this type

[Add a six-segment resource description to restrict the access](#)

Condition ☐ Source IP ⓘ [Add other conditions](#)

+ Add Permissions

Next Characters: 302 (up to 6,144)

Effect: Allow

Service: TencentDB for SQL Server (sqlserver)

Action: All actions

Resource: Specific resources > Add a six-segment resource description

Fill in the following items respectively: Resource prefix: instanceid and Resource ID: specific instance ID, as well as Resource prefix: user and Resource ID: specific username.

Note:

Method of determining the resource prefix: In the CAM-supported interfaces for TencentDB for SQL Server, there is a corresponding six-segment resource description. For details, refer to [TencentDB for SQL Server CAM-supported interfaces](#).

2. Click **Next**, name your custom policy, and assign the policy to the target sub-account.
3. Click **Complete** to finish the authorization.

Step 2: Enable CAM verification

There are two scenarios for enabling CAM verification: enabling CAM verification when creating an account and enabling CAM verification for an existing account. You can follow the steps below for each scenario.

Scenario 1: Enabling CAM verification when creating an account

Scenario 2: Enabling CAM verification for an existing account

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Instance ID** or click **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, choose **Account Management** > **Create Account**, enter relevant information in the pop-up window, and click **OK** after confirmation.

Note:

For detailed steps on creating an account with different permissions, refer to [Account Management](#). The following describes the steps related to enabling CAM verification.


Enable CAM verification: Turn on the switch for "Enable CAM verification", read the important notice in the pop-up window, and click **OK**.

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Instance ID** or click **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, choose **Account Management**.
4. On the account management page, find the target account and click **Enable CAM Verification** in its operations column.
5. Read the important notice in the pop-up window, and then click **OK**.


Step 3: Obtain the password through code calling in the application

Once the account has the relevant CAM permission specifications and CAM verification is enabled, you can obtain the password through code calling in Java or other languages in the application to connect to the database instance.

1. In the Tencent Cloud console, query the APPID of the account on the [Account Information](#) page.




2. Obtain the SecretID and SecretKey in [CAM Console > API Key Management](#).
3. Use the following code in the application.




```
<dependency>  
<groupId>com.tencentcloudapi</groupId>  
<artifactId>tencentcloud-dbauth-sdk-java</artifactId>  
<version>1.0.3</version>  
</dependency>
```

Indirect dependency: tencentcloud-sdk-java 3.1.1039 or later versions.



```
<dependency>
  <groupId>com.tencentcloudapi</groupId>
  <artifactId>tencentcloud-sdk-java</artifactId>
  <version>3.1.1039</version>
</dependency>
```

Example of obtaining the password through code calling



```
package com.tencentcloud.dbauth;
import com.tencentcloudapi.common.Credential;
import com.tencentcloud.dbauth.model.GenerateAuthenticationTokenRequest;
import com.tencentcloudapi.common.exception.TencentCloudSDKException;
import com.tencentcloudapi.common.profile.ClientProfile;
import com.tencentcloudapi.common.profile.HttpProfile;

public class GenerateDBAuthentication {

    public static void main(String[] args) {
        // Define the parameters for an authentication token.
    }
}
```

```
String region = "<Instance region>";
String instanceId = "<Instance ID>";
String userName = "<Username>";
// Get the credentials from an environment variable.
Credential credential = new Credential(System.getenv("<TENCENTCLOUD_SECRET_

System.out.println(getAuthToken(region, instanceId, userName, credential));
}
```

```
public static String getAuthToken(String region, String instanceId, String user
try {
    // Instantiate an HTTP profile, which is optional and can be skipped if
    HttpProfile httpProfile = new HttpProfile();
    httpProfile.setEndpoint("cam.tencentcloudapi.com");
    // Instantiate a client profile, which is optional and can be skipped i
    ClientProfile clientProfile = new ClientProfile();
    clientProfile.setHttpProfile(httpProfile);

    // Build a GenerateAuthenticationTokenRequest.
    GenerateAuthenticationTokenRequest tokenRequest = GenerateAuthenticatio
        .region(region)
        .credential(credential)
        .userName(userName)
        .instanceId(instanceId)
        .clientProfile(clientProfile) // clientProfile is optional.
        .build();

    return DBAuthentication.generateAuthenticationToken(tokenRequest);

} catch (TencentCloudSDKException e) {
    e.printStackTrace();
}
return "";
}
```

<Instance ID>: Replace with the ID of the instance you need to access.

<TENCENTCLOUD_SECRET_ID>: Replace with the SecretID obtained from the CAM console.

<TENCENTCLOUD SECRET KEY>: Replace with the SecretKey obtained from the CAM console.


Step 4: Use the identity token to connect to TencentDB for SQL Server

Note:

Using the JDBC driver for connection is the standard way for Java programs to connect to relational databases.

Detailed installation and connection methods for the JDBC driver can be found in [Using the JDBC Driver](#).

After obtaining the identity token AuthToken in [Step 3](#), you can use it to connect to TencentDB for SQL Server. The following connection command is an example for connecting to the database using JDBC.



```
String connectionString = "jdbc:sqlserver://localhost:1433;databaseName=<Database name>";
Connection con = DriverManager.getConnection(connectionString);
```

<Database name>: Replace with the name of the database you need to access.

<Username>: Replace with the actual username to log in.

<Password>: Replace with the AuthToken obtained in [Step 3](#).

Appendix 1: Resetting Password

When the CAM verification feature is enabled for the account, you can update the password through the password reset operation. If the account is set to change the password every 12 hours, you can immediately update the password before the rotation cycle is reached by performing the password reset operation.

Note:

Note that the current login credentials will become invalid after the password is reset. You need to check whether the database access status meets expectations.

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Instance ID** or click **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, choose **Account Management**.
4. On the account management page, find the target account and click **Reset Password** in its operations column.
5. Read the risk warning in the pop-up window, and then click **OK**.

Appendix 2: Disabling CAM Verification

Note:

After disabling CAM verification, you will not be able to obtain access credentials through CAM. Please update your password promptly.

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Instance ID** or click **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, choose **Account Management**.
4. On the account management page, find the target account and click **Disable CAM Verification** in its operations column.
5. In the pop-up window, enter the new password and confirm the password, and then click **OK**.

Database Management

Database List

Last updated : 2024-01-18 17:20:33

This document describes the fields in the account list.

Database account


After you create a TencentDB for SQL Server instance, you can create different database accounts to allocate and manage databases based on your business needs.

Viewing the account list

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click **Manage** in the **Operation** column of the target instance to enter the **Instance Details** page.
3. Click **Account Management** to enter the **Account Management** tab.

Instance Details	System Monitoring	Backup and Restoration	Security Group	Account Management	Database Management	Bac
<div>Create Account</div> <div>Batch Management</div>						
<input type="checkbox"/>	Account Name	Account Type	Status	Account Creation T...	Account Update Ti...	Password Update T...
<input type="checkbox"/>	tor	Standard Account	Normal	2023-02-07 17:20:05	2023-02-07 17:20:46	2023-02-07 17:20:46

Tools

Tool	Description
Search box	Click  to filter accounts. You can fuzzy search for accounts by account name.
Refresh	Click


to refresh the list.

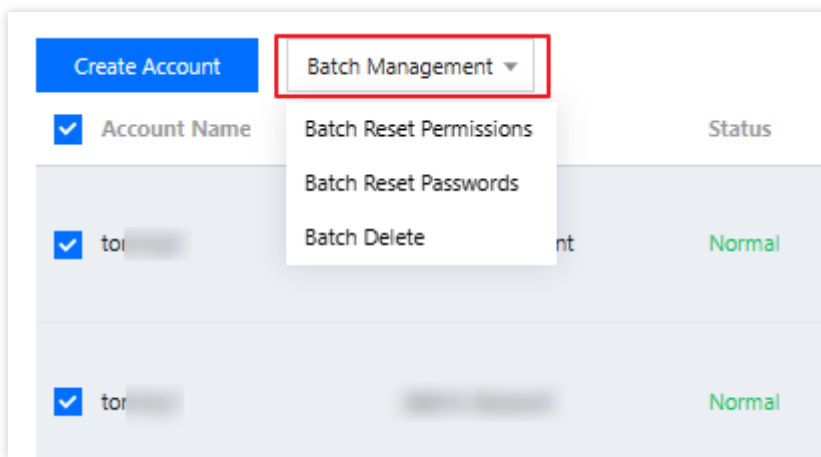
Fields in the account list

Field	Description
Account Name	The name of the account created for the instance.
Account Type	<p>The types of accounts created in the instance. The types of accounts supported by different instance architectures are as follows:</p> <p>Two-node</p> <p>Privileged account: Instance admin account, which has the owner permissions of all databases by default.</p> <p>Standard account: Non-admin account, which can be granted the read/write, read-only, or owner permissions of a specific database.</p> <p>Designated account: It can only view and own the specified database. A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.</p> <p>Single-node</p> <p>Admin account: Instance admin account, which has the highest-level sysadmin permission and the owner permissions of all databases. After the admin account is enabled, the product SLA will no longer be guaranteed.</p> <p>Privileged account: Instance admin account, which has the owner permissions of all databases by default.</p> <p>Standard account: Non-admin account, which can be granted the read/write, read-only, or owner permissions of a specific database.</p> <p>Designated account: It can only view and own the specified database. A designated account can be authorized to multiple databases, but a database can be authorized to only one designated account.</p>
Status	The status of the account.
Account Creation Time	The time when the account was created, which can be sorted in ascending or descending order.
Account Update Time	The time when the account was last updated, which can be sorted in ascending or descending order.
Password Update	The time when the account password was last reset, which can be sorted in ascending or descending order.

Time	
Database	The databases that the account is authorized to access.
Remarks	The remarks of the account.
Operation	Set Permissions: Set the read/write or read-only permission of the specified database for the account. Reset Password: Reset the password of the account. Delete Account: Delete the account.

Note:

For each operation in the account list, you can select multiple accounts and click the corresponding batch operation button above the list for batch management.



More

[Creating Account](#)[Modifying Account Permissions](#)[Resetting Instance Password](#)[Deleting Account](#)

Creating Database

Last updated : 2024-04-11 09:43:24

Overview

TencentDB for SQL Server allows you to create databases on the **Database Management** page in the console and authorize accounts for database access.

Maximum number of databases created in a single instance

Note:

In a High Availability or Cluster Edition instance, if you use the default value of 0 for the `max worker threads` parameter, you cannot create more than 100 databases. To create more databases, you must set this parameter to 20,000 as instructed in [Setting Instance Parameters](#).

If the instance only has one CPU core, we recommend that you keep the database limit at 70 to guarantee instance stability.

SQL Server 2008 R2 Enterprise instances don't support lifting the database quantity limit, which is 70. The limit in other SQL Server instances is subject to the number of instance CPU cores as calculated below:

Dual-Node (Formerly High-Availability Edition)

2012 Standard/Enterprise

2014 Standard/Enterprise

2016 Standard/Enterprise

Maximum number of databases:

$$\min\{80 + \sqrt{\text{CPU core quantity}} * 40, 300\} \text{ (Extract the square root and round it to one decimal place)}$$

Extract the square root of the CPU core quantity, round it to one decimal place, multiply the result by 40, and add the product to 80 to get the value X. The smaller value between X and 300 is the maximum number of databases. For example, you can create up to 160 databases in a 4-core 16 GB MEM SQL Server 2014 Enterprise instance.

Dual-Node (Formerly Cluster Edition)

2017 Enterprise

2019 Enterprise

2022 Enterprise

The maximum number of databases is related to the instance CPU. The maximum number of databases for instances with 8-core CPU or less is 80, and the maximum number of databases for instances with 8-core CPU or above is 100.

Single Node (Formerly Basic Edition)

2008 R2 Enterprise

2012 Enterprise

2014 Enterprise

2016 Enterprise

2017 Enterprise

2019 Enterprise

2022 Enterprise

Maximum number of databases:

$$\min\left\{\left\lfloor\sqrt{\text{CPU core quantity}}\right\rfloor * 100, 400\right\}$$

Extract the square root of the CPU core quantity, round it down to the nearest integer, and multiply the result by 100 to get the value N. The smaller value between N and 400 is the maximum number of databases. For example, you can create up to 200 databases in a 4-core 16 GB MEM SQL Server 2017 Enterprise instance.

Table of instance CPU core quantity and corresponding maximum database quantity

Maximum number of databases in Dual-Server High Availability/Cluster Edition

Maximum number of databases in Basic Edition

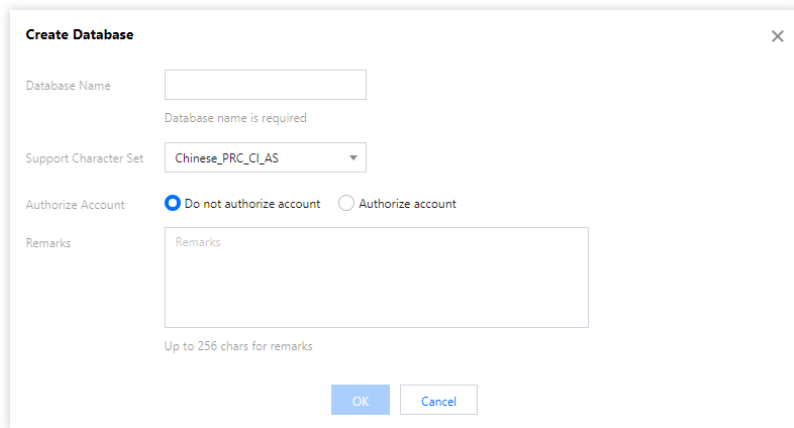
CPU Cores	Dual-Node 2008 R2/2012/2014/2016 Enterprise Edition	Dual-node 2017/2019/2022 Enterprise Edition
1	70	80
2	136	80
4	160	80
8	193	100
12	218	100
16	240	100
24	275	100
32	300	100
48	300	100

64	300	100
96	300	100

Number of CPU Cores	Single-Node 2008 R2 Enterprise Edition	Single-Node 2012/2014/2016/2017/2019/2022 Enterprise Edition
2	80	100
4	80	200
8	80	200
12	80	300
16	80	400
24	80	400
32	80	400
48	80	400
64	80	400

Directions

1. Log in to the [TencentDB for SQL Server console](#) and click the ID of the target instance in the instance list or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Database Management** tab, click **Create Database**, set configuration items in the pop-up window, confirm that everything is correct, and click **OK**.



The image shows a 'Create Database' dialog box with the following fields and options:

- Database Name:** A text input field with a placeholder 'Database name is required' below it.
- Support Character Set:** A dropdown menu currently showing 'Chinese_PRC_CI_AS'.
- Authorize Account:** Two radio buttons: 'Do not authorize account' (selected) and 'Authorize account'.
- Remarks:** A large text area with a placeholder 'Remarks' and a note 'Up to 256 chars for remarks' below it.
- Buttons:** 'OK' and 'Cancel' buttons at the bottom right.

Database Name: It can contain up to 32 letters, digits, and underscores and must start with a letter.

Supported Character Set: Select the character set to be used by the database. Currently, most native character sets are supported.

Authorize Account: You can authorize existing accounts to access the database. If you haven't created an account yet, see [Creating Account](#).

Remarks: It can contain up to 256 characters.

Setting Database Permissions

Last updated : 2024-01-18 17:20:34

Scenario

TencentDB for SQL Server allows you to grant and modify database account permissions. You cannot modify the permissions of admin and privileged accounts because they have owner permissions of all databases in a TencentDB for SQL Server instance. On the database management page, you can grant other accounts read/write, read-only, or owner permissions of created databases.


Prerequisites

You have created an account other than admin or privileged accounts in your TencentDB for SQL Server instance as instructed in [Creating Account](#).

You have created at least one database in your TencentDB for SQL Server instance as instructed in [Creating Database](#).

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select **Database Management**, find the target database, and click **Set Permissions** in its **Operation** column.

Instance Details	System Monitoring	Backup and Restoration	Publish/Subscribe	Security Group	Account Management	Database Management	Read-Only Instance	Backup Management
Slow Query Log	Parameter Configuration							
<div>Create DatabaseBatch Management</div>								
<input type="checkbox"/>	Database Name	Status	Character Set	Creation Time	Bind Account	Description	Operation	
<input type="checkbox"/>		Running		2021-11-16 15:05:49	qqq Read-Only	---	<div>Set PermissionsDeleteCloneOthers</div>	

4. In the pop-up window, select the target account on the left and permissions to be granted on the right and click **OK**.

Note:

You can batch set permissions. On the database management page, select multiple databases and click **Batch Management > Batch Reset Permissions** at the top.

Batch resetting permissions will clear all set database account permissions; that is, account permissions of the selected databases will be reset.

Cloning Database

Last updated : 2024-01-18 17:20:33

Overview

TencentDB for SQL Server provides the database cloning feature for you to quickly clone an existing database to your current instance. You need to specify the new database name during cloning, while other information items such as account permissions of the new database are the same as those of the source database.

Prerequisites

The source database is in **Running** status.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.

Instance Details	System Monitoring	Backup and Restoration	Publish/Subscribe	Security Group	Account Management	Database Management	Read-Only Instance	Backup Management
Slow Query Log	Parameter Configuration							
Create Database		Batch Management ▾						
<input type="checkbox"/>	Database Name	Status	Character Set	Creation Time	Bind Account	Description	Operation	
<input type="checkbox"/>		Running	Chinese_PRC_CI_AS	2021-11-16 15:05:49	qqq Read-Only	--	Set Permissions Delete Clone	
							Others ▾	

3. On the instance management page, select **Database Management**, find the target database, and click **Clone** in its **Operation** column.

Instance Details	System Monitoring	Backup and Restoration	Publish/Subscribe	Security Group	Account Management	Database Management	Read-Only Instance	Backup Management
Slow Query Log	Parameter Configuration							
Create Database		Batch Management ▾						
<input type="checkbox"/>	Database Name	Status	Character Set	Creation Time	Bind Account	Description	Operation	
<input type="checkbox"/>		Running	Chinese_PRC_CI_AS	2021-11-16 15:05:49	qqq Read-Only	--	Set Permissions Delete Clone	
							Others ▾	

4. In the pop-up window, name the new database and click **OK**.

Clone Database ×

Original Database Name	New Database Name
te	<input type="text" value="ne"/> ✓

OK Cancel

Deleting Database

Last updated : 2024-01-18 17:20:33

Overview

TencentDB for SQL Server supports deleting databases on the **Manage Database** tab in the console.

Note:

Batch database deletion is not supported currently.

Directions

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Manage Database** tab, select the database to be deleted in the list, and click **Delete** in the **Operation** column.

Note:

Before deleting the database, please make sure that there are no active connections using the database. If any active connection is detected during the deletion process, the deletion operation will be suspended.

Create Database

Batch Management

<input type="checkbox"/>	Database Name	Status	Character Set	Creation Time	Bind Account
<input type="checkbox"/>	thesf	Running	Chinese_PRC_CI_AS	2021-03-25 19:52:39	--

3. In the pop-up window, confirm the information and click **Delete**.

Setting Change Data Capture (CDC)

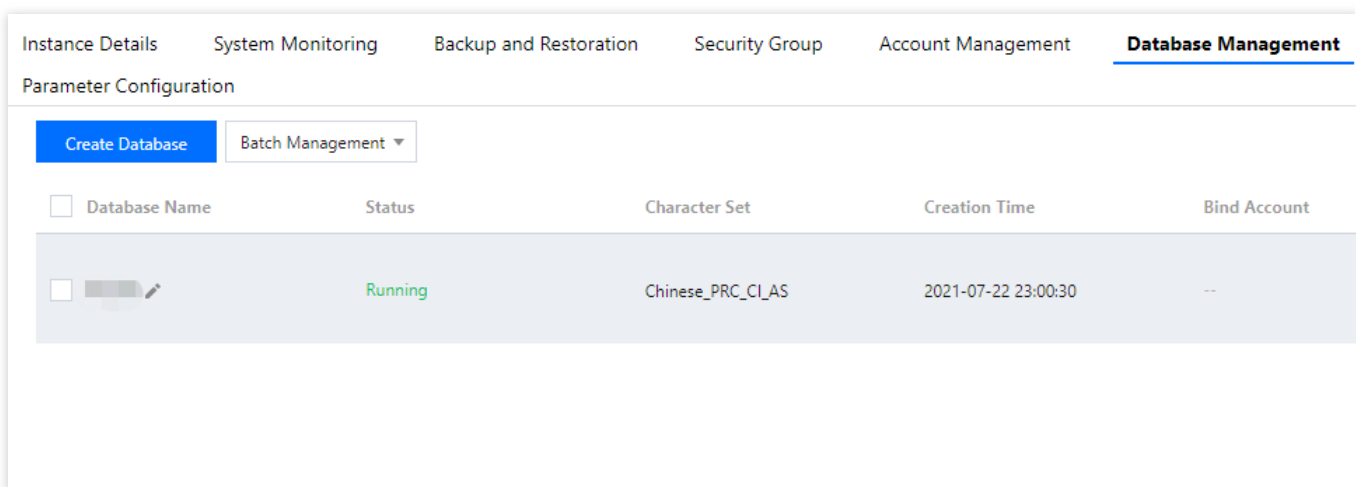
Last updated : 2024-01-18 17:20:33

Change data capture (CDC) is used to capture insertions, updates, and deletion applied to SQL Server tables and provide the details of these changes in a convenient relational format.

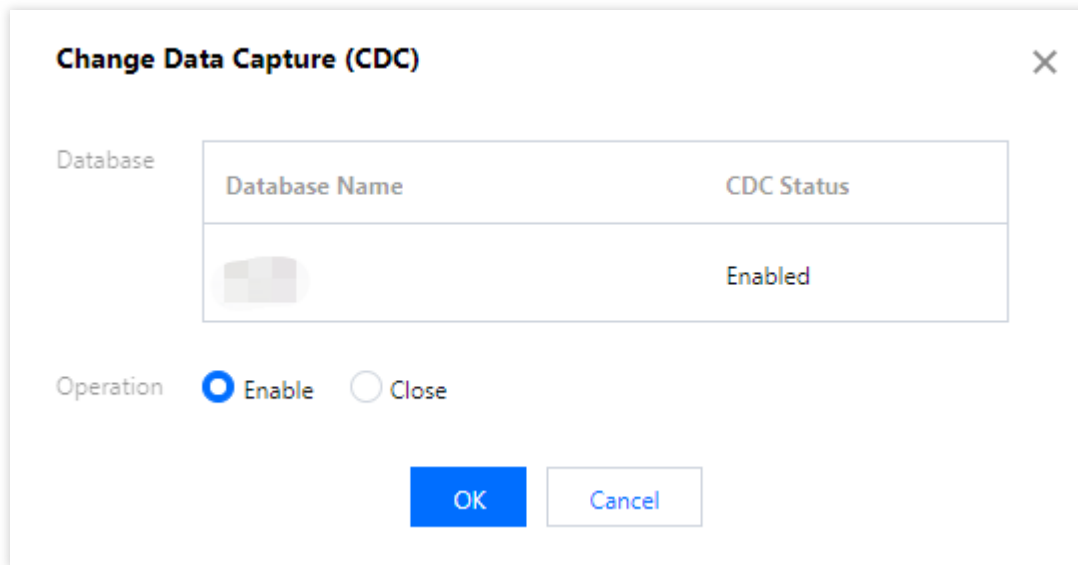
The change table used for CDC contains the columns of the column structure of the source table tracked by the image, as well as the metadata required for understanding the changes that have occurred. After CDC is enabled for a table, all DML and DDL operations on it will be recorded, which helps track changes to it.

Enabling/Disabling CDC for one database


1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the **Database Management** tab, locate the row of the target database and click **Other > Enable/Disable CDC** in the **Operation** column.



3. The pop-up window displays the name and current CDC status of the database. After enabling or disabling CDC, click **OK**.



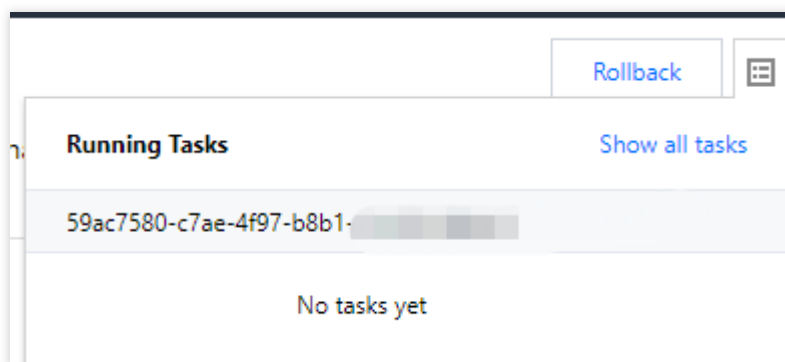
The dialog box is titled "Change Data Capture (CDC)" and has a close button (X) in the top right corner. It contains a "Database" section with a table showing the CDC status for a specific database. The table has two columns: "Database Name" and "CDC Status". The "Database Name" column contains a placeholder icon, and the "CDC Status" column contains the text "Enabled". Below the table, there is an "Operation" section with two radio buttons: "Enable" (selected) and "Close". At the bottom, there are two buttons: "OK" and "Cancel".

Database Name	CDC Status
	Enabled


Operation ☒ Enable ☐ Close

OK Cancel


You can view the task progress of enabling or disabling CDC through **Running Tasks** in the top-right corner of the **Database Management** tab.



The "Running Tasks" panel is located in the top-right corner of the "Database Management" tab. It has a "Rollback" button and a "Show all tasks" link. Below the header, there is a list of tasks. The first task is identified by the ID "59ac7580-c7ae-4f97-b8b1-" and has a progress bar. Below the list, it says "No tasks yet".

Rollback 

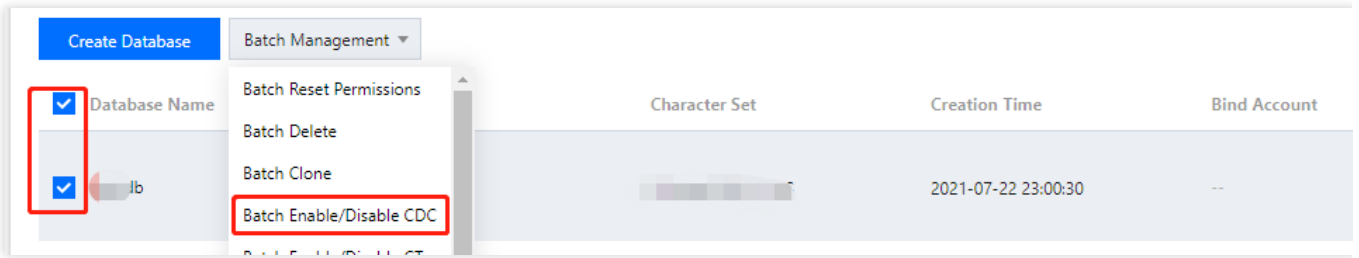
Running Tasks [Show all tasks](#)

59ac7580-c7ae-4f97-b8b1- 

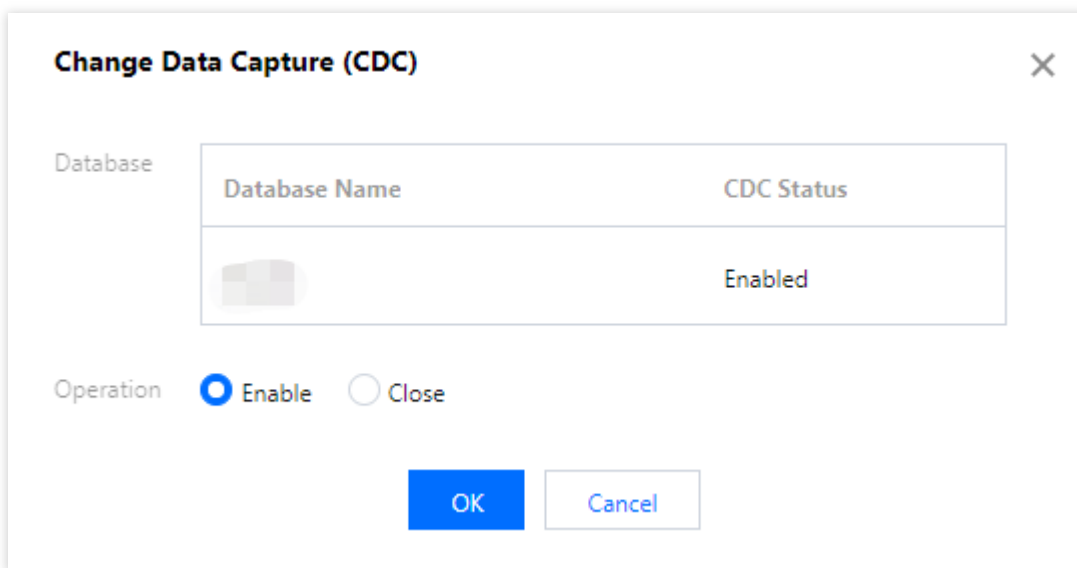
No tasks yet

Batch enabling/disabling CDC for multiple databases

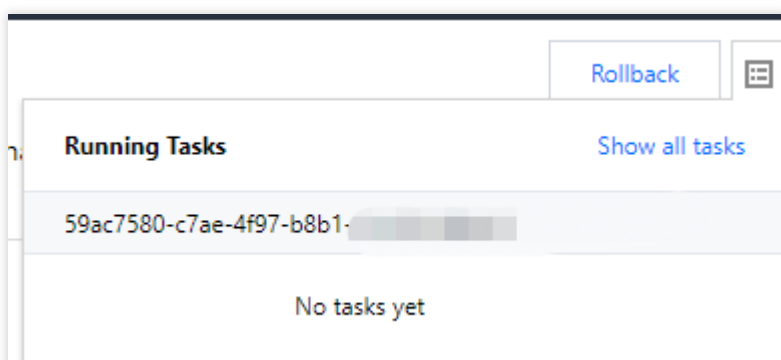
1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID to enter the instance management page.
2. On the instance management page, select the **Database Management** tab, select the rows of the target databases, and click **Batch Management** > **Batch Enable/Disable CDC** at the top of the list.



3. The pop-up window displays the names and current CDC status of the selected databases. After enabling or disabling CDC, click **OK**.



You can view the task progress of enabling or disabling CDC through **Running Tasks** in the top-right corner of the **Database Management** tab.



Setting Change Tracking (CT)


Last updated : 2024-01-18 17:20:33

Change tracking (CT) can be applied to track a specific table or even column in a database. When you perform an addition, modification, or deletion operation in a table with CT enabled, the system will automatically generate a version number for the operation and record the operation information, including the operation timestamp, operation type, and primary key of the affected data.

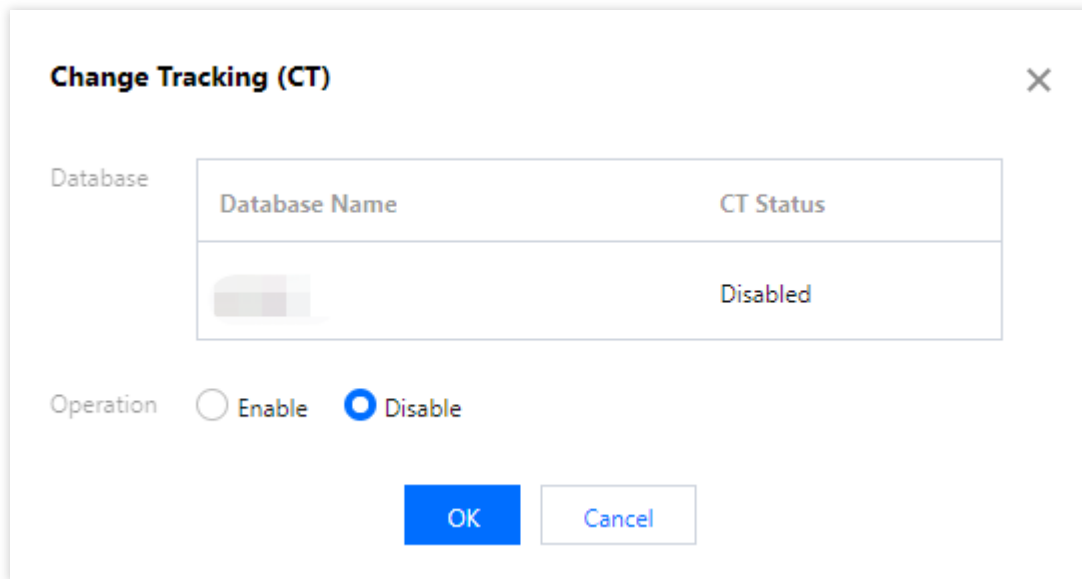
CT only records whether the rows in the table have been changed but not the changed data, so it has little impact on the performance of data operations. Such information will be recorded in the SQL Server system table and automatically cleared and maintained by the system.

Enabling/Disabling CT for one database


1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the **Database Management** tab, locate the row of the target database and click **Other** > **Enable/Disable CT** in the **Operation** column.



3. The pop-up window displays the names and current CT status of the databases. After enabling or disabling CT, click **OK**. If you enable CT, you can set the data retention period.




You can view the task progress of enabling or disabling CT through **Running Tasks** in the top-right corner of the **Database Management** tab.

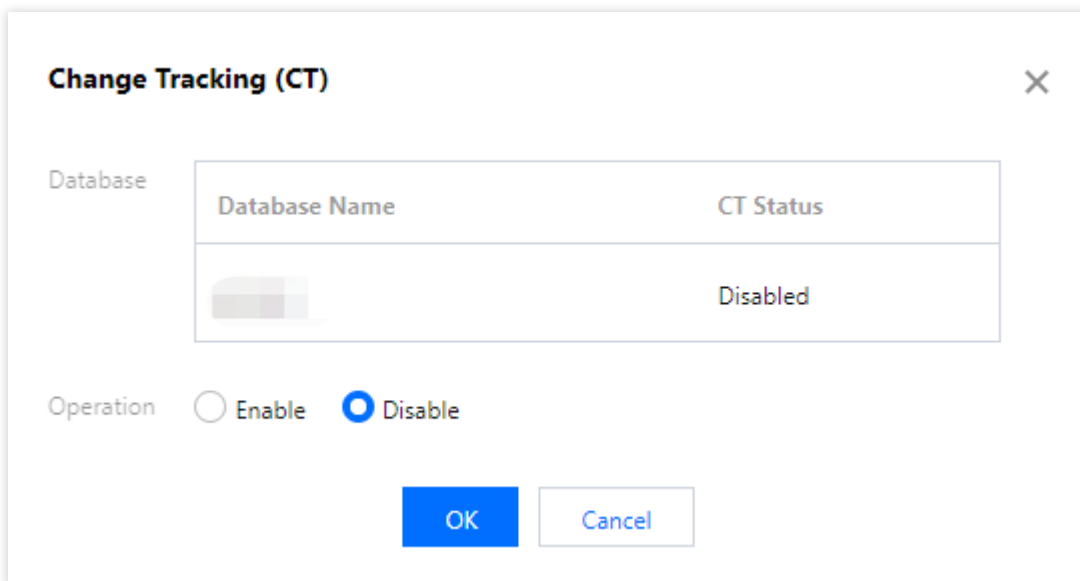


Batch enabling/disabling CT for multiple databases

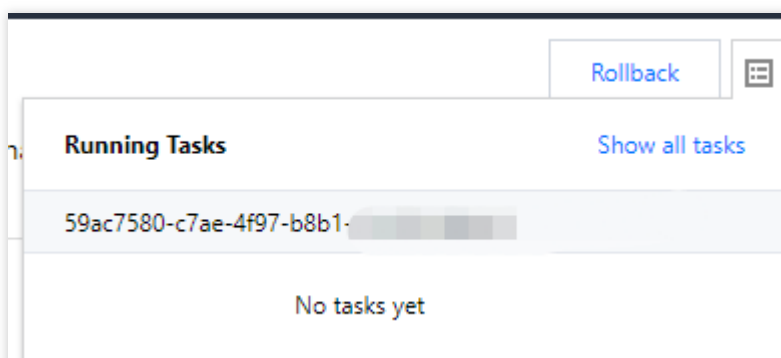
1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID to enter the instance management page.
2. On the instance management page, select the **Database Management** tab, select the rows of the target databases, and click **Batch Management** > **Batch Enable/Disable CT** at the top of the list.



3. The pop-up window displays the names and current CT status of the databases. After enabling or disabling CT, click **OK**. If you enable CT, you can set the data retention period.



You can view the task progress of enabling or disabling CT through **Running Tasks** in the top-right corner of the **Database Management** tab.



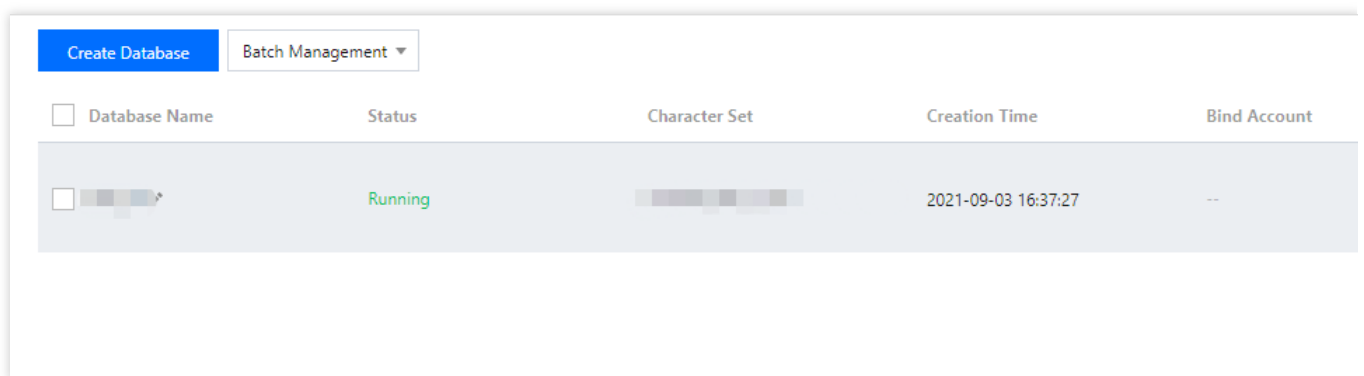
Shrinking Database

Last updated : 2024-01-18 17:20:33

TencentDB for SQL Server supports reducing database size by shrinking all database files to release unused space. This document describes how to shrink a database in the TencentDB for SQL Server console.

Shrinking one database

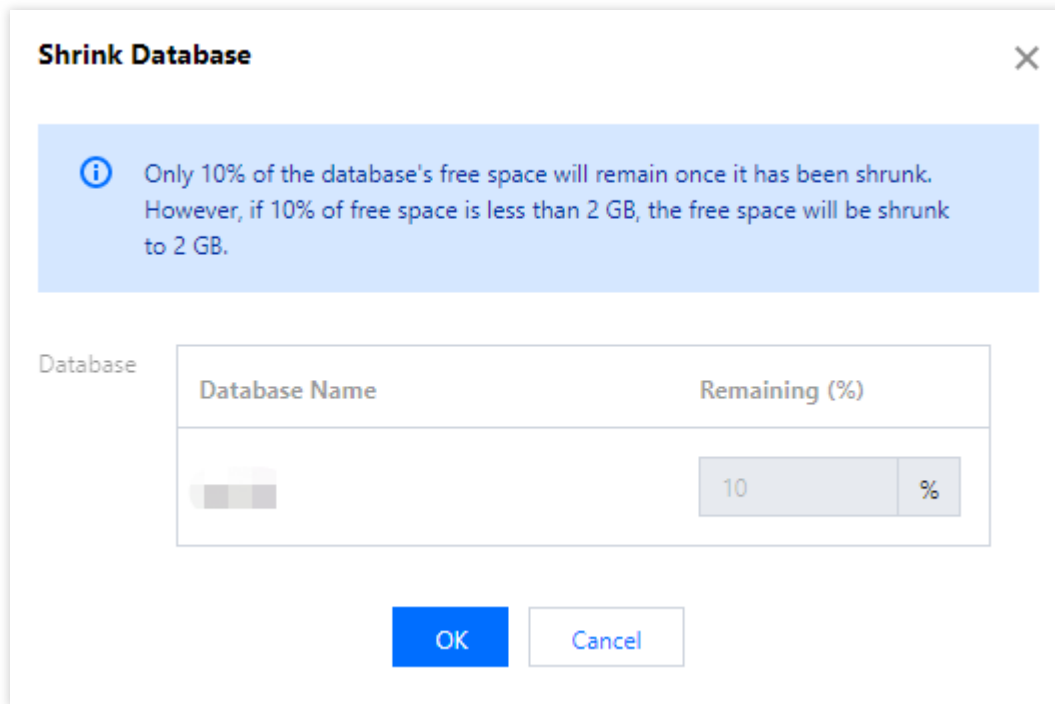
1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the **Database Management** tab, locate the row of the database to be shrunk and click **Other > Shrink Database** in the **Operation** column.




3. The pop-up window displays the database name and the ratio of remaining space. Currently, only shrinking to 10% of the free space is supported.

Note:

Only 10% of the database's free space will remain once it has been shrunk. However, if 10% of free space is less than 2 GB, the free space will be shrunk to 2 GB.

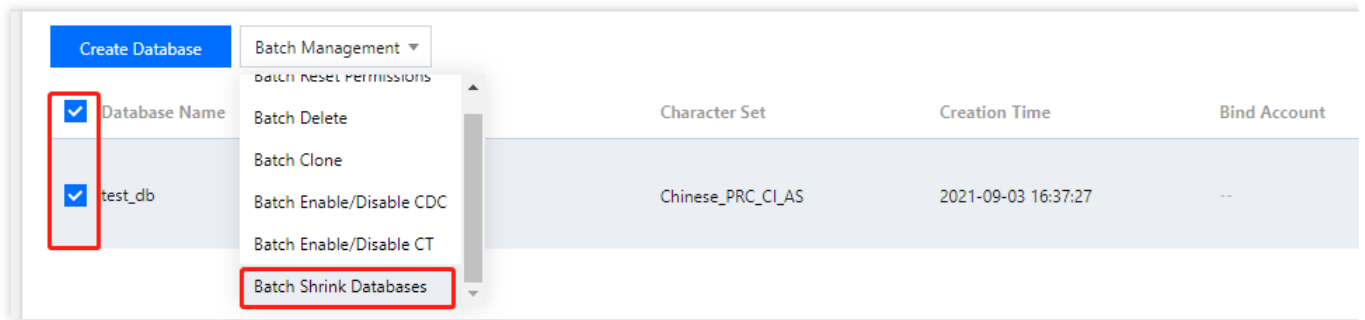


You can view the task progress of database shrinking through **Running Tasks** in the top-right corner of the **Database Management** tab.



Batch shrinking databases

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the **Database Management** tab, select the rows of the databases to be shrunk and click **Batch Management** > **Batch Shrink Databases** at the top of the list.



3. The pop-up window displays the database name and the ratio of remaining space. Currently, only shrinking to 10% of the free space is supported.

Note:

Only 10% of the database's free space will remain once it has been shrunk. However, if 10% of free space is less than 2 GB, the free space will be shrunk to 2 GB.

Shrink Database



Only 10% of the database's free space will remain once it has been shrunk. However, if 10% of free space is less than 2 GB, the free space will be shrunk to 2 GB.

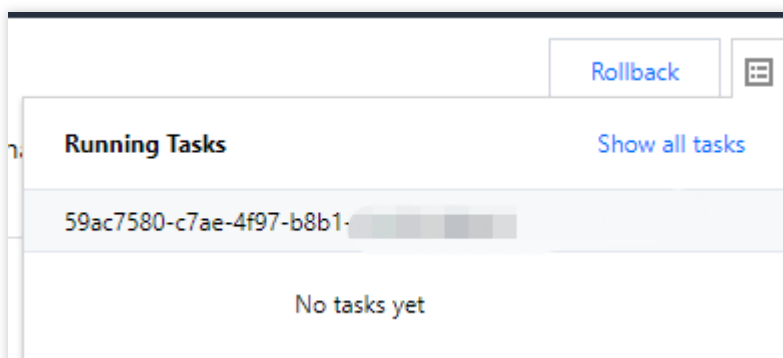
数据库

Database Name	Remaining (%)
	<input type="text" value="10"/> %
	<input type="text" value="10"/> %
	<input type="text" value="10"/> %

OK

Cancel

You can view the task progress of database shrinking through **Running Tasks** in the top-right corner of the **Database Management** tab.



Data Security

Transparent data encryption

Last updated : 2024-01-18 17:23:30

This document introduces the Transparent Data Encryption feature of TencentDB for SQL Server.

Feature Overview

Transparent Data Encryption (TDE) is a feature that allows for the encryption and decryption of data to be transparent to the user. TDE provides file-level encryption, supporting real-time I/O encryption and decryption of data files. It can be transparent to applications at the database layer, requiring no modification of business code. Encryption is performed before data is written to the disk, ensuring that data stored on the disk is encrypted. Decryption occurs when data is read from the disk into memory. It means that encryption and decryption are transparently performed during disk data read and write operations and TDE does not increase the size of data files, which meets the compliance requirements for static data encryption.

Use Cases

TDE is usually used to address security and compliance issues in various scenarios, where the static data needs to be protected, such as PCI DSS and CCP compliance.

Functional Limitations

Only supports Tencent Cloud's automatically generated key certificates, user-generated key certificates are not supported.

Keys and certificates cannot be downloaded.

TDE at the instance level can only be enabled and cannot be disabled. Once TDE is enabled, the certificates and keys on the instance cannot be deleted.

TDE at the database level can be enabled or disabled.

All instances under the same account (UIN) have the same encryption certificate, which is used for backup recovery and rollback between different instances.

For instances under different accounts (UIN), if backup recovery and rollback are needed, the encryption certificate under the same account must be referenced.

Business intelligence service instances do not support the enablement of TDE.

Functional Description

The encryption certificate of TDE is an account-level certificate, and each account can only have one certificate, and each account can only have one TDE certificate, indicating that different instances under the same account have the same encryption certificate. This allows you to perform backup restoration and rollback among instances without disabling TDE.

For instances that reference the same encryption certificate, during backup recovery and data migration, if TDE is not enabled for the target instance but is enabled for the source database, then it will be automatically enabled for the target instance without affecting your business, as the certificates are the same under the same account.

If TencentDB for SQL Server database instances are distributed across different Tencent Cloud accounts and require cross-account backup and recovery, select Reference Certificate from Other Accounts as the certificate source when enabling TDE. By doing so, you can ensure that the certificates are the same for the database instances with TDE enabled under different accounts. That is, a root account (UIN) can only have one certificate, and one certificate may be used by multiple root accounts (UIN).

If a master instance is associated with RO instance, you only need to enable data encryption for the instance, and the data encryption for the RO instance will be enabled automatically.

If the instance is associated with a read-only instance or publish/subscribe, disassociate them before enabling or disabling instance TDE at the **database level**.

After TDE is enabled, data cannot be restored by a backup file offline. To restore it to a local database, you need to disable TDE and create a manual backup to restore data.

Databases encrypted offline cannot be directly migrated to TencentDB for SQL Server instances. You need to first disable the TDE function at the offline database level before migrating.

TDE enhances the data security while compromising the read/write performance of encrypted databases and significantly increasing CPU utilization. Therefore, enable TDE with caution. In addition, it is not recommended to enable the TDE function for instances with less than 4 CPU cores.

It may take a long time to enable or disable TDE, during which any operation on the instance is not supported, otherwise, you may fail to enable or disable TDE. We recommend that you perform this operation during the off-peak hours. These tasks include but are not limited to:

- Modifying, deleting database, making database offline, detaching database.

- Converting database or file group to read-only status.

- Backing up database.

- Rolling back/restoring database.

- Changing data capture (CDC) .

- Changing tracking (CT) .

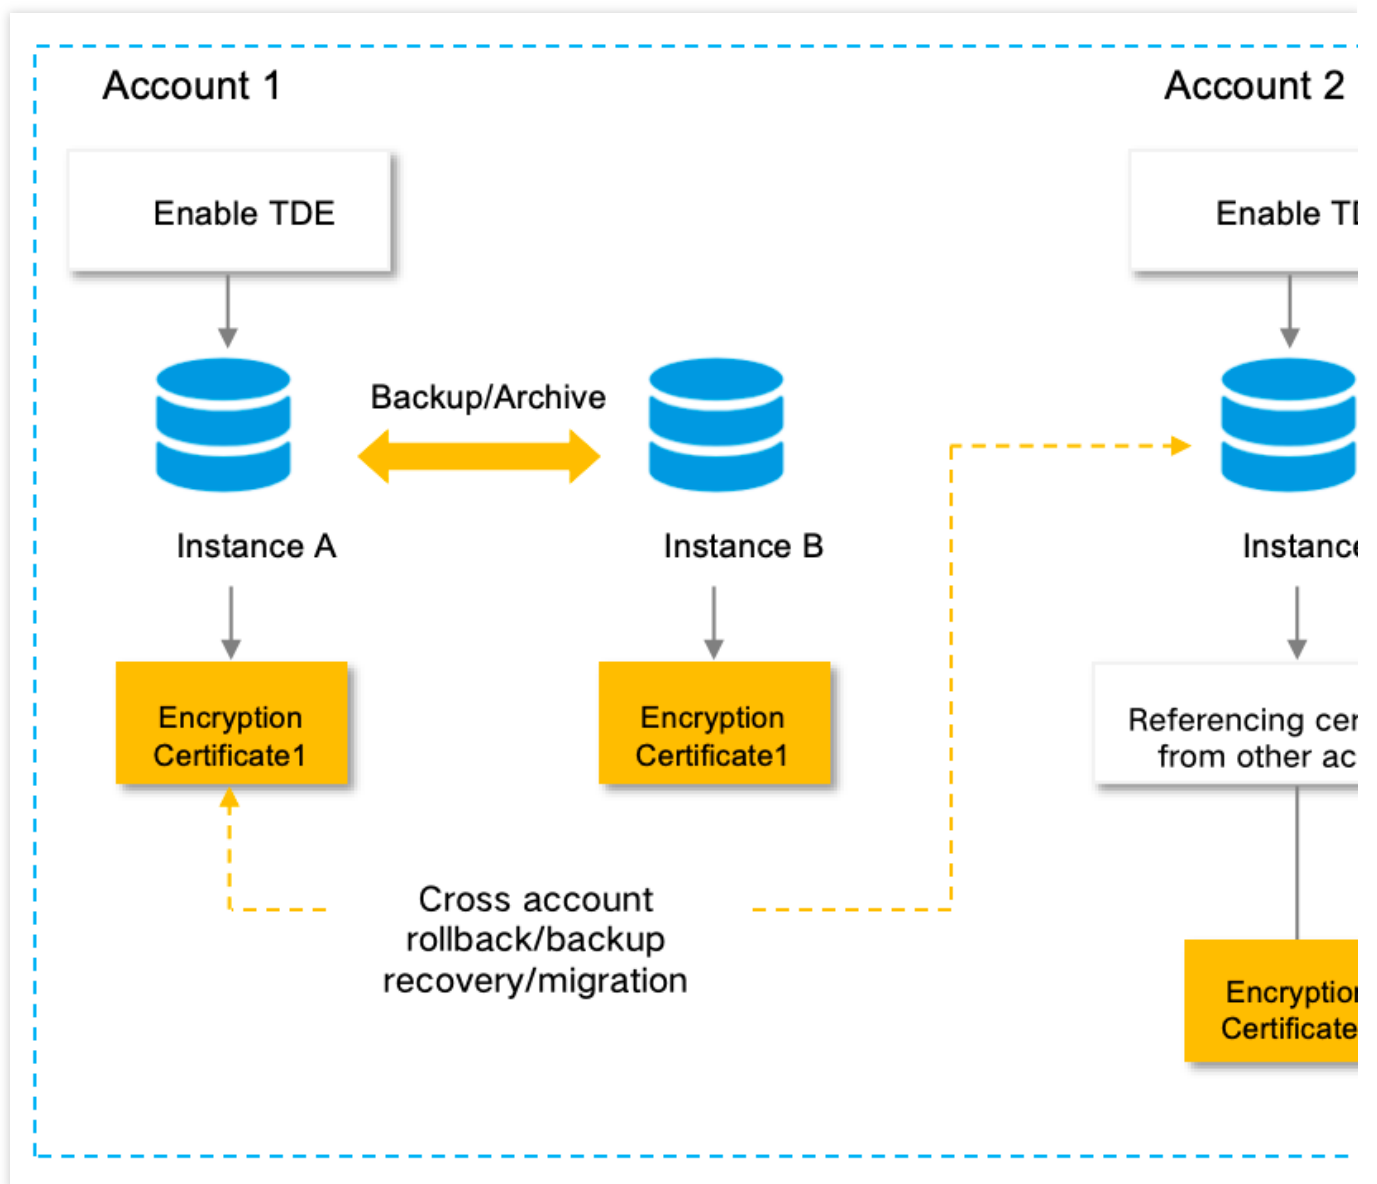
- Shrinking database.

- Cloning database.

- Modifying database permissions.

Note:

If users possess multiple root accounts (UIN), all of which require the enablement of the TDE feature, and subsequent cross-account backup recovery/migration/rollback operations are necessary among different root accounts (UIN), it is imperative to ensure that all root accounts reference the same certificate. That is, after the instance in Account 1 has enabled the TDE feature, the instances in Account 2, Account 3, Account 4, and so forth, must unequivocally reference the encryption certificate of Account 1 when enabling the TDE feature. Otherwise, subsequent cross-account operations (including but not limited to: cross-account rollback, cross-account backup recovery, cross-account migration, etc.) will be unfeasible.



Enabling Instance Transparent Data Encryption

Note:

After enabling Transparent Data Encryption, the encryption certificate is generated by Tencent Cloud, and other encryption certificates not provided by Tencent Cloud cannot be used.

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region, then in the instance list, click an **Instance ID** or **Manage** in the **Operation** column to enable TDE.
3. In the **instance management** page, select **Data Security** > **Data Encryption**, then click the **button** to enable the feature, located next to the TDE status.

Note:

The instance-level TDE feature cannot be disabled once enabled.

Instance Details System Monitoring Backup and Restoration Publish/Subscribe Security Group Account Management Database Management

Parameter Configuration **Data Security**

Data Encryption

- Each account can only have one TDE certificate, indicating that different instances under the same account have the same encryption certificate. This allows you to perform backup restoration and recovery.
- During backup recovery and data migration, if TDE is not enabled for the target instance but is enabled for the source database, then it will be automatically enabled for the target instance without the need to manually enable it.
- If the database instances are distributed across different Tencent Cloud accounts and require cross-account backup and recovery, select "Reference Certificate from Other Accounts" as the certificate.
- After TDE is enabled, data cannot be restored by a backup file offline. To restore it to a local database, you need to disable TDE and create a manual backup to restore data.
- TDE enhances the data security while compromising the read-write performance of encrypted databases and significantly increasing CPU utilization. Therefore, enable TDE with caution.
- It may take a long time to enable or disable TDE, during which any operation on the instance is not supported, otherwise, you may fail to enable or disable TDE. We recommend that you perform this operation during off-peak hours.
- If the instance is associated with a read-only instance or the pub/sub paradigm, disassociate them before enabling or disabling instance TDE for encryption or decryption.
- If a primary instance is associated with RO instance, you only need to enable data encryption for the instance, and the data encryption for the RO instance will be enabled automatically.

TDE Status ☐ Disabled (cannot be disabled once enabled)

4. In the pop-up TDE encryption dialog box, there exist three scenarios. You may operate according to the actual scenario.

Scenario One: No Encryption Certificate

Enable TDE



Once enabled, this feature will not be disabled. For more information, see [Enabling Transparent Data Encryption](#).

Certificate Source



Use Certificate of This Account



Reference Certificate from Other Accounts

If you use the certificate of this account and have enabled TDE for a database, you will not be able to restore data to a database instance under other account. To restore data to a cross-account instance, disable TDE on the source database and create a manual backup to restore data.

OK

Cancel

Select the source of the certificate, with options to **Use Certificate of This Account** or **Reference Certificate from Other Accounts**.

Use Certificate of This Account

Given that the encryption certificate for TDE is an account-level certificate, it implies that this account has not previously enabled TDE for any instance under it. Select to **Use Certificate of This Account** and click **OK**.

Reference Certificate from Other Accounts



This denotes the use of encryption certificates from other accounts. Select the option to **Reference Certificate from Other Accounts**, choose Reference Account, and then click **OK**.

Note:

If you reference certificate from other accounts, the certificates are the same for the instances with TDE enabled under different accounts. Therefore, cross-account backup and recovery can be smoothly performed, which means that backup files can be used to restore data to database instances under other accounts.

Scenario Two: Use Certificate of This Account

Enable TDE

 Once enabled, this feature will not be disabled. For more information, see [Enabling Transparent Data Encryption](#) .

Certificate Source ☒ Use Certificate of This Account
☐ Reference Certificate from Other Accounts

If you use the certificate of this account and have enabled TDE for a database, you will not be able to restore data to a database instance under other account. To restore data to a cross-account instance, disable TDE on the source database and create a manual backup to restore data.

OK

Cancel

This scenario signifies that the account has previously enabled TDE for a certain instance under the account, and select the certificate source as Use Certificate of This Account. In such circumstances, the certificate source will default to **Use Certificate of This Account**. Simply clicking **OK** will enable TDE.

Scenario Three: Reference Certificate from Other Accounts

This scenario signifies that the account has previously enabled TDE for a certain instance under the account, with the certificate source Reference Certificate from Other Accounts. In such circumstances, the certificate source will default to **Reference Certificate from Other Accounts**. After selecting the Reference Account, you can click **OK** to enable TDE.

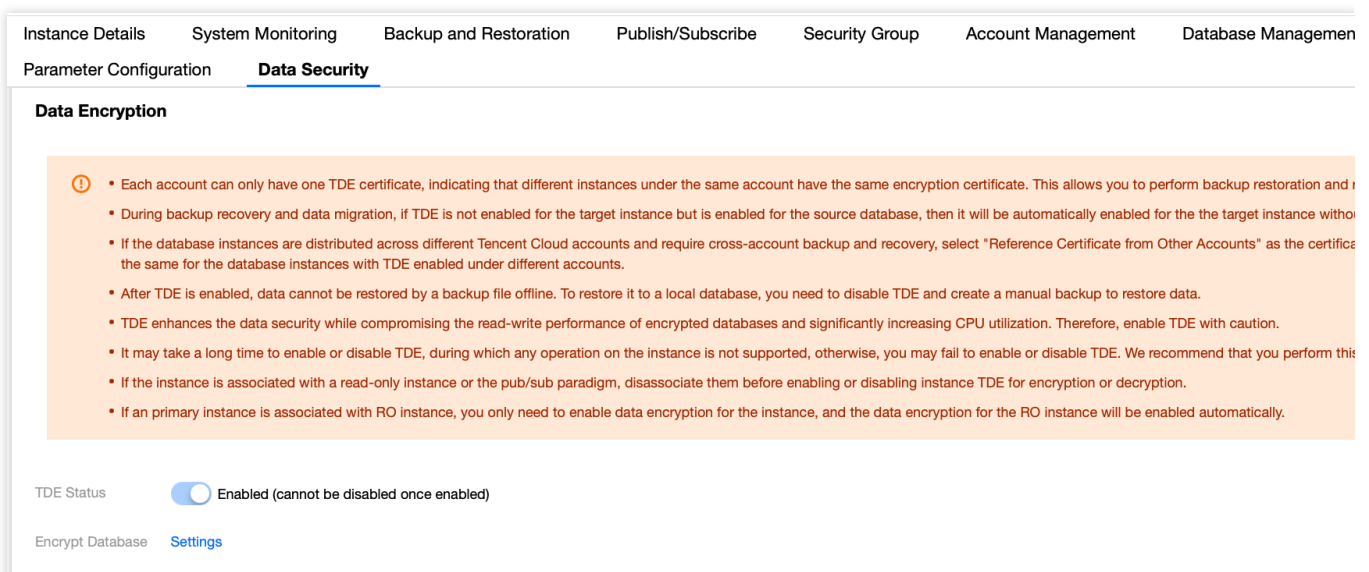
Enabling or Disabling Encrypted Databases

Note:

To enable/disable the TDE feature at the database level for an instance, the instance must not be associated with any read-only instances or publish/subscribe relationships. First, disassociate any read-only instances or publish/subscribe relationships, then enable/disable the TDE feature at the database level, and finally re-add the read-only instances and publish/subscribe relationships.

The prerequisite for setting up an encrypted database is that the Transparent Data Encryption feature of the instance has already been enabled. For the procedure, please refer to [Enabling Instance Transparent Data Encryption](#).

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region, then in the instance list, click an **Instance ID** or **Manage** in the **Operation** column to enable TDE.
3. In the **Instance Management** page, select **Data Security** > **Data Encryption**, and then click on **Settings** after encrypting the database.



4. In the pop-up window, select the desired databases from the Unencrypted Database on the left, indicating the enablement of encryption for these databases. Conversely, remove databases from the Encrypted Database on the right, signifying the disablement of encryption for these databases. Click **OK** after performing operations as required.

Database Encryption

Unencrypted Database(0)

Enter the database name

☐ Database Name

Encrypted Database (0)

Enter the database name

Database Name

↔

OK

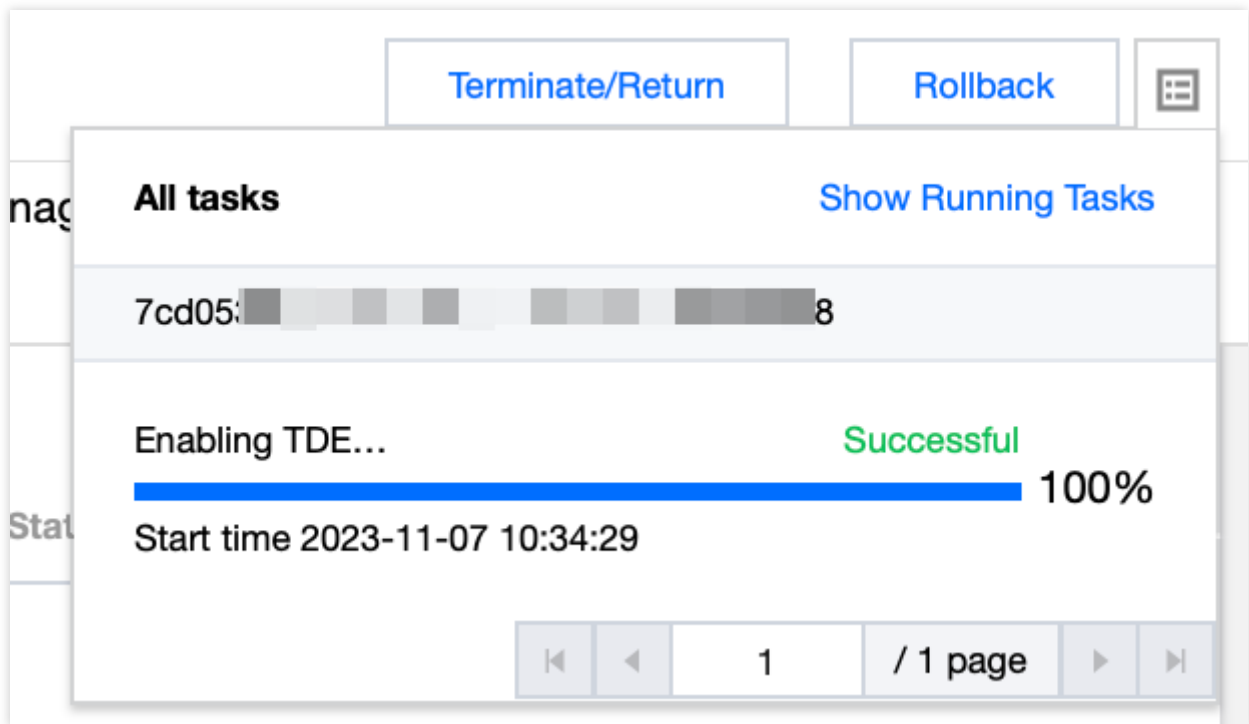
Cancel

5. After enabling the TDE function for a specific database dimension, one can view the databases with encryption functions either enabled or disabled under database management, based on the **TDE status** field.

Instance Details	System Monitoring	Backup and Restoration	Publish/Subscribe	Security Group	Account Management	Database Management
Parameter Configuration	Data Security					
Create Database	Batch Management					
<input type="checkbox"/> Database Name	Status	Character Set	Creation Time	Bind Account	Descripti	

Reviewing Tasks

When you enable or disable the instance-level TDE feature or the database-level TDE feature, you can understand the current task progress through the task icon in the upper right corner of the console.



Constraints on Cloud Databases with TDE Feature Enabled during Rollback, Backup Recovery, Migration, and Database Cloning

Feature Set	Specific Features	Description
Rollback	Revert to Source Instance	Consistent with the instance certificate, a rollback can be performed.
	Revert to Other Existing Instances under the Same Account	Certificates initialization is consistent across different instances under the same account. During the rollback process, the system will determine whether the conditions of the source database encryption being disabled or the target instance encryption being enabled are met before proceeding with the rollback.
	Cross-Regional Rollback	Enable instances with cross-regional backup, synchronizing encrypted cross-regional backup files. Utilizing encrypted cross-regional backups for offsite rollback requires the certificates of instances in different regions under the same account to be uniformly initialized. During the rollback, the system will determine whether the source database has encryption disabled

		or the target instance has encryption enabled before proceeding with the rollback.
Clone Database	Clone to Source Instance	Consistent with the instance certificate, cloning can be performed.
Backup Restoration (Cold Migration)	Restoration to Source Instance from Backup	Consistent with the instance certificate, backup and restoration can be performed.
	Backup Restoration to Other Existing Instances under the Same Account	When restoring from an encrypted backup file, it is necessary to ensure that the source backup file's encryption has been disabled or the target instance's encryption has been enabled. Therefore, if the source database has enabled TDE encryption and the target instance has not initiated instance-level TDE encryption, the system will automatically enable the instance-level TDE function for the target instance, given that the encryption certificates under the same account are identical.
	Restoration of Backup to Cross-Account Instances	When restoring from an encrypted backup file, it is necessary to ensure that the source backup file's encryption has been disabled or the target instance's encryption has been enabled. Therefore, if the source database has enabled TDE encryption and the target instance has not initiated instance-level TDE encryption, the system will automatically enable instance-level TDE functionality for the target instance if the source database and target instance under the same account reference the same certificate. However, if the source database and target instance under different accounts reference different certificates, it is required to either disable the encryption of the source database or enable the encryption of the target instance and reference the same certificate before proceeding with the backup restoration.
Migration with DTS	DTS Migration to Source Instance	Consistent with the instance certificate, migration can be conducted.
	DTS Migration to Other Existing Instances under the Same Account	When migrating encrypted files, it is necessary to ensure that the source backup file encryption has been disabled or the target instance encryption has been enabled. Therefore, if the source database has enabled TDE encryption and the target instance has not initiated instance-level TDE encryption, the system will automatically enable the instance-level TDE function for the target instance, given that the encryption certificates under the same account are identical.
	DTS Migration to Cross-Account	Encrypted files, during data migration, necessitate the condition of either the source backup file having encryption disabled or the

	Instances	target instance having encryption enabled to proceed with data migration. Consequently, when the source database has TDE encryption enabled and the target instance does not have instance-level TDE encryption enabled, if the source database and the target instance under the same account reference the same certificate, the system will automatically enable instance-level TDE functionality for the target instance. However, if the source database and the target instance under different accounts reference different certificates, it is required to either disable encryption on the source database or enable encryption on the target instance and reference the same certificate before proceeding with data migration.
Publish/Subscribe	Link from Source Instance to Other Existing Instances under the Same Region	When publishing and subscribing, the system will determine that the source database encryption must be disable or the target instance encryption must be enable before publishing and subscribing can be performed.

Access Management

CAM Overview

Last updated : 2024-01-18 17:23:30

Issues

If you use Tencent Cloud services, including CVM, VPC, and TencentDB, which are managed by different people who share your Tencent Cloud account key, the following issues may arise:

There is a high risk of key leakage because the key is shared among multiple individuals.

The absence of limitations on other users' access rights may easily lead to incorrect operations, causing security risks.

Solutions

You can avoid the issues above by providing different users with sub-accounts, permitting them to manage different services. By default, a sub-account does not possess the authorization to utilize Tencent Cloud services or the related resources. Consequently, we should formulate a policy to allow sub-accounts to use the resources and permissions they need.

[Cloud Access Management \(CAM\)](#) aids in the secure and convenient management of access to Tencent Cloud services and resources. With CAM, you can create sub-accounts, user groups, and roles, controlling their access scope through a policy. CAM supports SSO capabilities for users and roles, allowing targeted settings for interaction between corporate users and Tencent Cloud based on specific management circumstances. Your initially created Tencent Cloud root account possesses complete access to all Tencent Cloud services and resources. It is recommended to safeguard your root account credentials, utilize sub-accounts or roles for daily access, enable multi-factor authentication, and change keys regularly.

While CAM is used, a policy can be associated with a user or a group of users to allow or reject the use of specific resources by users to accomplish designated tasks. For more information on CAM policies, please refer to [Policy Syntax](#).

If you do not need to manage the CAM of the related resources of the Tencent Cloud Database for the sub-accounts, you may bypass this part. It will not impede your comprehension or usage of the remaining parts in this document.

Quick Start

A CAM policy must authorize or deny the use of one or more cloud database operations. Simultaneously, it must specify the resources that can be used for these operations, which could be all the resources (some operations can also be partial resources). The policy can also encompass the conditions stipulated for the operated resources.

Note :

Users are recommended to use CAM policies to manage TencentDB resources and authorize TencentDB operations. While the experience for existing users with project-based permissions remains unchanged, it is not suggested to continue resource management and operation authorization with project-based permissions. The TencentDB does not support the setting of related validity conditions for the time being.


Task	Link
Understanding the fundamental structure of policies	Policy Syntax
Defining operations in a policy	Operations in TencentDB
Defining resources in a policy	Resources Path in TencentDB
Resource-level permissions supported by TencentDB	Resource-level authorization supported by TencentDB

Authorization Policy Syntax

Last updated : 2024-01-18 17:23:30

Policy Syntax

CAM Policy:



```
{
```

```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "effect",
      "action": ["action"],
      "resource": ["resource"],
      "condition": {"key": {"value": {}}}
    }
  ]
}
```

Version: This field must be filled in; currently only the value 2.0 is acceptable.

Statement: Describes the detailed information of one or multiple permissions. It comprises permissions or collections of permissions for multiple other elements like effect, action, resource, and condition. Each policy contains just one statement element.

Effect: This field must be filled in. It describes whether the outcome of a statement is Allow or Explicitly Deny. The outcome only includes these two scenarios.

Action: This field must be filled in. It is used to describe the operation of Allow or Deny. An operation can be a API, which is prefixed with **sqlserver:**.

Resource: This field must be filled in. It describes the specific data of authorization. The resource is described in a six-segment format. Detailed resource outlines can vary with different products.

Condition: This field must be filled in. It describes the conditions under which the policy comes into effect. The conditions include an operator, an action key, and an action value. Condition values encompass time and IP addresses. Certain services also permit users to specify different values within these conditions.

SQL Server Operations

In the SQL Server policy statement, you can specify any API operation from any service supporting SQL Server. APIs prefixed with **sqlserver:** should be used for SQL Server, such as **sqlserver:DescribeDBInstances** or **sqlserver:CreateAccount**.

To specify multiple operations within a single statement, please separate them with a comma as demonstrated below:



```
"action":["sqlserver:action1","sqlserver:action2"]
```

You may also use an asterisk wildcard to specify multiple operations. For instance, you can designate all the operations with the name beginning with **Describe**, as shown below:



```
"action":["sqlserver:Describe*"]
```

To specify all the operations in SQL Server, please use an asterisk wildcard (*), as indicated below:



```
"action":["sqlserver:*"]
```

SQL Server Resources

Each CAM policy statement has its own resources.

The typical format of resources is as follows:



```
qcs:project_id:service_type:region:account:resource
```

project_id: Describes the project information, which is only used to enable compatibility with legacy CAM logic and can be left empty.


service_type: The product's abbreviation, such as sqlserver.

region: Describes the regional information, such as ap-guangzhou.

account: The root account information of the resource owner, such as uin/65xxx763.

resource: Indicates the detailed resource information of each product, such as instance/instance_id1 or instance/*.

For instance, you may use the specific instance (mssql-m8oh024t) to specify a resource in the statement as demonstrated below:



```
"resource": [ "qcs::sqlserver:ap-guangzhou:uin/65xxx763:instance/mssql-m8oh024t" ]
```

You could also employ an asterisk wildcard (*) to designate all instances pertaining to a certain account, as shown below:



```
"resource": [ "qcs::sqlserver:ap-guangzhou:uin/65xxx763:instance/*"]
```

If you want to specify all the resources or if a specific API operation does not support resource-level permissions, you can utilize an asterisk wildcard (*) within the resource element as shown below:



```
"resource": ["*"]
```

To specify multiple resources concurrently within a single command, segregate them with commas. The example of designation of two resources are as follows:



```
"resource":["resource1","resource2"]
```

The table below describes the resources that can be utilized by SQL Server and their corresponding description methods. In this context, words prefixed with \$ are considered placeholders. **Region** refers to a geographical area.

Account signifies the account ID.

Resources	Resource Description Method in Authorization Policies
Instances	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code>
VPC	<code>qcs::vpc:\$region:\$account:vpc/\$vpcId</code>

DFW	qcs::cvm:\$region:\$account:sg/\$sgId
-----	---------------------------------------

Authorizable Resource Types

Last updated : 2024-01-18 17:23:30

Resource-level permissions refer to the ability to specify which resources users are allowed to perform operations on. Some operations of SQL Server support resource-level permissions, meaning you can control when a user is allowed to perform operations or what specific resources they are permitted to use. The types of resources that can be authorized through Cloud Access Management (CAM) are as follows:

Resource Type	Resource Description Method in Authorization Policies
TencentDB instance-related	<code>qcs::sqlserver:\$region:\$account:instance/*</code> <code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code>

TencentDB for SQL Server supports resource-level authorization, allowing you to allocate specified sub-accounts with API permissions for specified resources. The following table presents TencentDB API operations currently supporting resource-level permissions, along with the supported resources and conditional keys for each operation. When specifying a resource path, you can utilize an * wildcard in the path.

Note :

Any TencentDB API operation not listed in the table does not support resource-level permissions. You can still authorize a user to perform these operations, but you must specify the * as the resource element in the policy statement. The table below showcases only a portion of the resource types. For more information, please refer to [Authorizable Resource Types for SQL Server](#).

API Name	API Description	Six-Segment Example of Resource
CreateAccount	Creating an account	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
CreateBackup	Creating a backup	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
CreateDB	Creating a database	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DeleteAccount	Deleting an account	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DeleteDB	Deleting a database	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeAccounts	Querying	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code>

	an account list	<code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeBackups	Querying a backup list	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeDatabaseNames	Querying a database name	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeDBInstances	Querying instance lists	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeDBs	Querying a database list	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeInstanceTasks	Querying instance tasks	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeRollbackTime	Querying the time range available for rollback	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
DescribeSlowlogs	Querying slow log lists	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
InquiryPriceRenewDBInstance	Querying the price of renewed instances	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
InquiryPriceUpgradeDBInstance	Querying the price of upgraded instances	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyAccountPrivilege	Modifying account permissions	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyAccountRemark	Modifying	<code>qcs::sqlserver:\$region:\$account:instance/\$instanceId</code>

	account remarks	<code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyBackupStrategy	Modifying the time for cold backup	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyDatabasePrivilege	Modifying database permissions	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyDBInstanceName	Renaming instances	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyDBInstanceProject	Modifying instance project	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyDBName	Renaming a database	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ModifyDBRemark	Modifying database remarks	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
RenewDBInstance	Renewing instances	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
ResetAccountPassword	Resetting account password	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
RestartDBInstance	Restarting instances	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
RestoreInstance	Restoring cold backup instances	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
RollbackInstance	Rolling back instances	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
TerminateDBInstance	Terminating instances	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code> <code>qcs::sqlserver:\$region:\$account:instance/*</code>
UpgradeDBInstance	Upgrading	<code>qcs::sqlserver:\$region:\$account:instance/\$ir</code>

	Instances	qcs::sqlserver:\$region:\$account:instance/*
--	-----------	--

Configuring Security Groups

Last updated : 2024-01-18 17:23:30

Overview

[Security group](#) serves as a stateful virtual firewall with filtering feature for configuring network access control for one or more TencentDB instances. It is an important network security isolation tool provided by Tencent Cloud. Instances with the same network security isolation demands in one region can be put into the same security group, which is a logical group. TencentDB and CVM share the security group list and are matched with each other within the security group based on rules. For specific rules and limitations, see [Security Group Overview](#).

Note:

TencentDB for SQL Server security groups currently only support network access control for VPCs and public network but not the classic network.

As TencentDB doesn't have any active outbound traffic, outbound rules don't apply to it.

TencentDB for SQL Server security group supports primary and read-only instances.

Configuring a security group for TencentDB

Step 1. Create a security group

1. Log in to the [CVM console](#).
2. Select **Security Group** on the left sidebar, select a region, and click **Create**.
3. In the pop-up window, set the following configuration items, confirm that everything is correct, and click **OK**.

Template: Select a template based on the service to be deployed on the TencentDB instance in the security group, which simplifies the security group rule configuration. The configuration is shown in the table below:

Template	Description	Applicable Scenario
Open all ports	All ports are opened to the public and private networks by default. This may pose security issues.	-
Open ports 22, 80, 443, and 3389 and the ICMP protocol	Ports 22, 80, 443, and 3389 and the ICMP protocol are opened to the public network by default. All ports are opened to the private network.	This template doesn't take effect for TencentDB.
Custom	You can create a security group and then add custom rules. For detailed directions, see "Step 2. Add a security group rule" below.	-

Name: Custom name of the security group.

Project: Select a project for easier management. By default, **DEFAULT PROJECT** is selected.

Remark: A short description of the security group for easier management.

Step 2. Add a security group rule

1. On the [Security Group](#) page, click **Modify rule** in the **Operation** column on the row of the security group for which to configure a rule.

2. On the security group rule page, click **Inbound rules > Add rule**.

3. In the pop-up window, set the rule.

Type: **Custom** is selected by default. You can also choose another system rule template. **SQL Server(1433)** is recommended.

Source: Traffic source (inbound rules) or target (outbound rules). You need to specify one of the following options:

Source or Target	Description
A single IPv4 address or an IPv4 range	In CIDR notation, such as <code>203.0.113.0</code> , <code>203.0.113.0/24</code> or <code>0.0.0.0/0</code> , where <code>0.0.0.0/0</code> indicates all IPv4 addresses will be matched.
A single IPv6 address or an IPv6 range	In CIDR notation, such as <code>FF05::B5</code> , <code>FF05:B5::/60</code> , <code>::/0</code> or <code>0::0/0</code> , where <code>::/0</code> or <code>0::0/0</code> indicates all IPv6 addresses will be matched.
ID of referenced security group. You can reference the ID of: Current security group Other security group	Current security group indicates the CVM associated with the security group. Other security group indicates the ID of another security group under the same project in the same region.
Reference an IP address object or IP address group object in a parameter template .	-

Protocol Port: Enter the protocol type and port range or reference a protocol/port or protocol/port group in a [parameter template](#).

Note:

To connect to TencentDB for SQL Server, port 1433 must be opened.

Policy: **Allow** or **Reject**. **Allow** is selected by default.

Allow: Traffic to this port is allowed.

Reject: Data packets will be discarded without any response.

Remark: A short description of the rule for easier management.

4. Click **Complete**.

Use cases

Scenario: You have created a TencentDB for SQL Server instance and want to access it from a CVM instance.

Solution: When adding security group rules, select SQL Server(1433) in "Type" to open port 1433.

You can also set **Source** to all or specific IPs (IP ranges) as needed to allow them to access TencentDB for SQL Server from a CVM instance.

Inbound or Outbound	Type	Source	Protocol and Port	Policy
Inbound	SQL Server: 1433	All IPs: 0.0.0.0/0 Specific IPs: Specify IPs or IP ranges	TCP:1433	Allow

Step 3. Configure a security group

A security group is an instance-level firewall provided by Tencent Cloud for controlling inbound traffic of TencentDB. You can associate a security group with an instance when purchasing it or later in the console. The operations for configuring security groups in two scenarios are as detailed below:

Note:

Currently, security groups can be configured only for **TencentDB for SQL Server instances in VPC**.

Scenario 1: Associate a security group with an instance when purchasing it

After the security group is created, you can associate a security group with an instance when purchasing it, and also you can quickly locate the target group by multiple selection and fuzzy search.

1. Log in to [TencentDB for SQL Server purchase page](#).
2. Click the parameter **Security Group>Existing Security Group**, and select the target security group in the box. Multiple selection and fuzzy search are supported for quickly locating the target group.

Security Group

Create Security Group

Existing Security Groups

sg-8... | Mong..., sg-q... | ..., 80, 443, 3... ^

Select up to 5 securi

Search by security group name/ID is supported.

Q

30912162

sg-8... | Mon

sg-q... | ..., 80, 443, 33

to TencentDB for SQL Ser

Instance Config

Instance Name

Set After Creation

Set Now

3. After setting all the parameters, click **Buy Now**.

Note:

You can delete the redundant associated security groups after selecting multiple of them. At least one security group is reserved by default.

Security Group

Create Security Group

Existing Security Groups

sg-89e... | Mong..., sg-... | ..., 80, 443, 3... v

Select up to 5 security groups, 1

sg-89... | Mon x

sg-q... | ..., 80, 443, 338... | ... 62241530 [

Ensure that the port 1433 is opened for security group, otherwise, you can't connect to TencentDB for SQL Server. You ca

View Security Group Rules v

Scenario 2: Associate a security group with an instance after purchasing it in the console

1. Log in to the [TencentDB for SQL Server Console](#). In the instance list, select the instance for which to configure a security group and click **Manage** in the "Operation" column to enter the instance management page.
2. Select **Security Group** tab, and click **Configure Security Group**.
3. In the pop-up dialog box, select the security group to be bound and click **OK**.

Importing security group rules

1. On the [Security Group](#) page, click the ID/name of the target security group.
2. On the inbound rule or outbound rule tab, click **Import rule**.
3. In the pop-up window, select an edited inbound/outbound rule template file and click **Import**.

Note:

As existing rules will be overwritten after importing, we recommend that you export the existing rules before importing new ones.

Cloning a security group

1. On the [Security Group](#) page, locate the desired security group and click **More > Clone** in the **Operation** column.
2. In the pop-up window, select the target region and project and click **OK**. If the new security group needs to be associated with a CVM instance, do so by managing the CVM instances in the security group.

Deleting a security group

1. On the [Security Group](#) page, find the security group to be deleted and click **More > Delete** in the **Operation** column.
2. In the pop-up window, click **OK**. If the current security group is associated with a CVM instance, it must be disassociated first before being deleted.

SSL Encryption Settings

Last updated : 2024-01-26 16:02:28

SSL encryption overview

Secure Sockets Layer (SSL) authentication signifies the process of authenticating the connection from the clients to the cloud server, applicable to both the users and server. Enabling SSL encryption allows you to acquire a CA certificate and upload it to the server. When the clients access the database, the SSL protocol is activated. An SSL secure channel between the clients and the cloud server is established. This mechanism ensures encrypted data transmission, which prevents data interception and unauthorized modification, thus ensuring the safe transmission of information from both parties.

The SSL protocol needs to be established based on reliable TCP and has the advantage of being independent from application layer protocols. Therefore, high-level application layer protocols such as HTTP, FTP, and TELNET can be transparently established based on it. It completes encryption algorithm processing, communication key negotiation, and server authentication before communication is made over application layer protocols. After that, all data transferred over application layer protocols will be encrypted to ensure communication privacy.

Background

SSL, a security confidentiality protocol proposed by Netscape, constructs secure channels for data transmission between browsers and Web servers. It employs encryption algorithms such as RC4, MD5, and RSA to actualize secure communication. The Internet Engineering Task Force (IETF) standardized SSL 3.0, renaming it as the Transport Layer Security (TLS) after standardization. However, since the term SSL is more commonly used, SSL encryption in this document actually refers to TLS encryption.

Note:

The versions of TLS supported by Tencent Cloud Database are 1.0, 1.1, and 1.2.

In scenarios where a database is connected in an unencrypted manner, all information being transmitted over the network is in plaintext, thus posing the risk of being intercepted, tampered with, and impersonated by illegal users.

The SSL protocols are designed to solve these risks and can theoretically achieve:

The information is propagated in an encrypted manner, impervious to interception by any third parties.

There is a verification mechanism for immediate tampering detection by both parties in the communication.

Identity certificates will be used to authenticate the identity.

TencentDB for SQL Server enhances linkage security through the employment of SSL encryption, supporting the download and installation of SSL CA certificates to the corresponding application services.

Note:

Not protecting the data itself, SSL encryption guarantees the security of the traffic between the clients and the cloud database server. By encrypting the network connection at the transport layer, it can boost the security and integrity of the communication data. But this will concurrently increase the network connection response time.

Supported Architectures and Versions

The database instances of all the architectures and versions of TencentDB for SQL Server support SSL encryption settings. In this settings, read-only instances do not need to be configured separately. Once SSL encryption is enabled on a primary instance, it will automatically apply to its read-only instances.

Note:

If you set SSL on the primary instance and choose to execute it during maintenance time, the read-only instances associated with the primary instances will follow the primary instances to take effect. The SSL will take effect along with a restart during the maintenance time of the primary instance.

Notes

After the SSL encryption is enabled, the CPU usage of the instances will rise and the delay in read-write operations will increase. It is suggested that SSL encryption should be enabled only when there is a demand for encryption in the public network linkage. The private network linkage is relatively safe and usually does not require encryption. Enabling SSL encryption, updating SSL Certificates provided by Tencent Cloud, or disabling SSL encryption will restart the database instance. The instance may be unavailable for a few minutes. Therefore, ensure that your business has a reconnection mechanism before operation. It is recommended that should be done during off-peak hours.

The SSL certificate has a validity period (one year). Please manually renew the validity period of the certificate in the console [prior to its expiration](#), otherwise, the client programs using SSL encrypted connections will not be able to connect normally.


If the SSL certificate is not renewed after expiration, it will only cause the client programs using encrypted connections to be unable to connect with the instance normally. But it will not affect the normal operation of the instance or data security.

After SSL encryption is disabled, the instance can only be connected through non-SSL encrypted methods.

Enabling SSL encryption

1. Log in to the [TencentDB for SQL Server console](#).

2. Select **Region**. Click **Instance ID** that needs to enable SSL in the instance list or **Manage** in the **Operation** column.
3. On the **Instance Management** page, select **Data Security > SSL Encryption**. Then, click the button to enable the feature after **SSL encryption status**.



4. Select the execution time in the pop-up window, and click **Confirm**.

Enable SSL

i Enabling SSL encryption enhances the security of the communication. After SSL encryption is enabled on the primary instance, the secondary and read-only instances will also have SSL encryption in effect simultaneously during the period when SSL encryption takes effect on the primary instance. For more information, see [SSL Encryption](#).

Certificate Source

 Tencent Cloud

Execution Time **i**

Immediate execution

During maintenance

Maintenance Time: 00:00-06:00 (modify on the "Instance Maintenance" page)

☒ Your database instance will be restarted when SSL encryption is enabled. Please ensure that your business has a recovery mechanism.

OK

Cancel

Certificate Origin: By default, the certificate provided by Tencent Cloud is used.

Execution time:

Immediate Execution: SSL encryption is activated immediately after you click **OK**.

During Maintenance Time: SSL encryption is activated during the instance maintenance time. [Modification of the instance maintenance time](#) can be done on the instance details page.

Note:

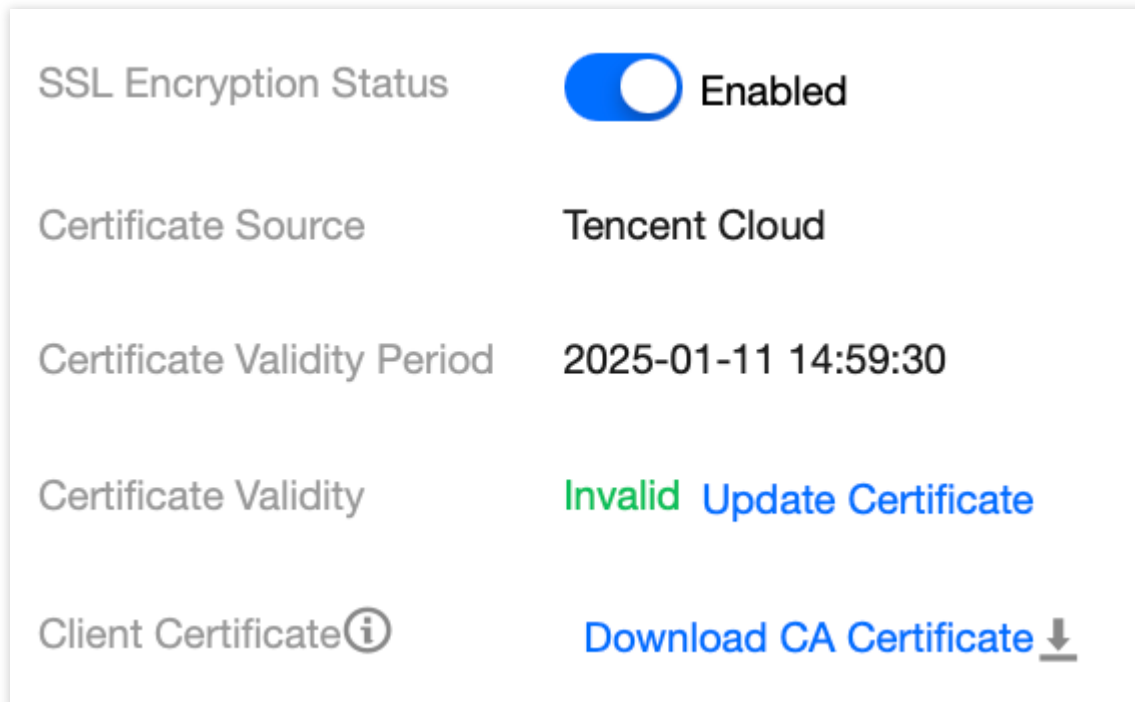
For detailed description of instance status changes during the activation of SSL encryption, please refer to [Appendix 1](#).

In the process of enabling SSL encryption, your database instance will be restarted. Be sure that your business has a reconnection mechanism.

5. Once

activated successfully

, the interface will appear as follows.



Please note that configuring the client CA certificate is an optional setting, which is used for the clients to trust the server. Click **Download CA Certificate** to download the certificate and install it onto the clients.

The downloaded file is a compressed package (TencentDB-SSL-CA.zip), containing the following three files:

.p7b file: It is used to import the CA certificate in Windows.

.jks file: It is a truststore certificate storage file in Java with a unified password **tencentdb**, which is used to import the CA certificate chain in the Java program.

.pem file: It is used to import the CA certificate in other systems or applications.

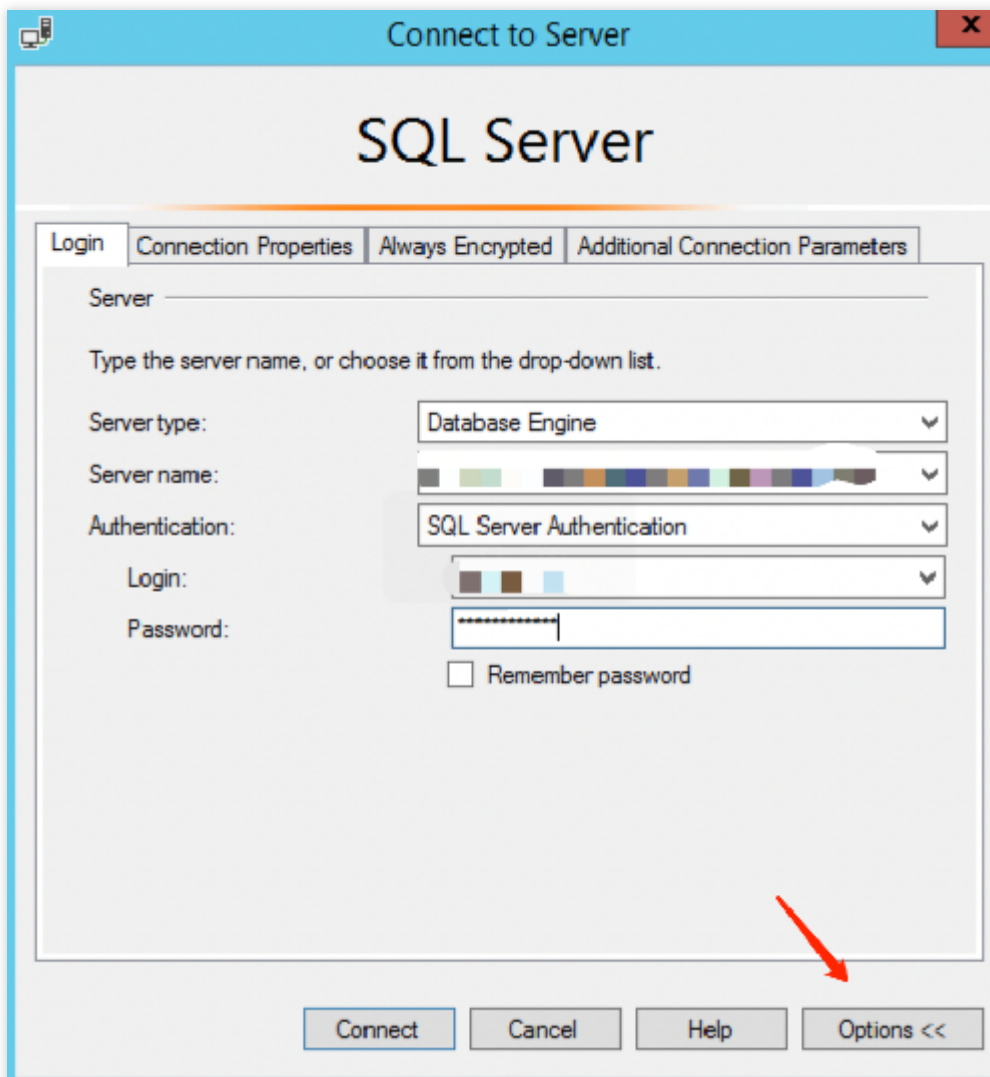
Connecting to SSL VPN Client

Once SSL encryption is enabled, when the clients are logging in to the TencentDB for SQL Server instances, there exist two scenarios of the trusted server certificate and the untrusted server certificate. If the server certificate is trusted, there's no need to configure SSL CA certificate for encrypted connections. However, if it isn't trusted, you are required to import and set up the SSL CA certificate for establishing the encrypted connection.

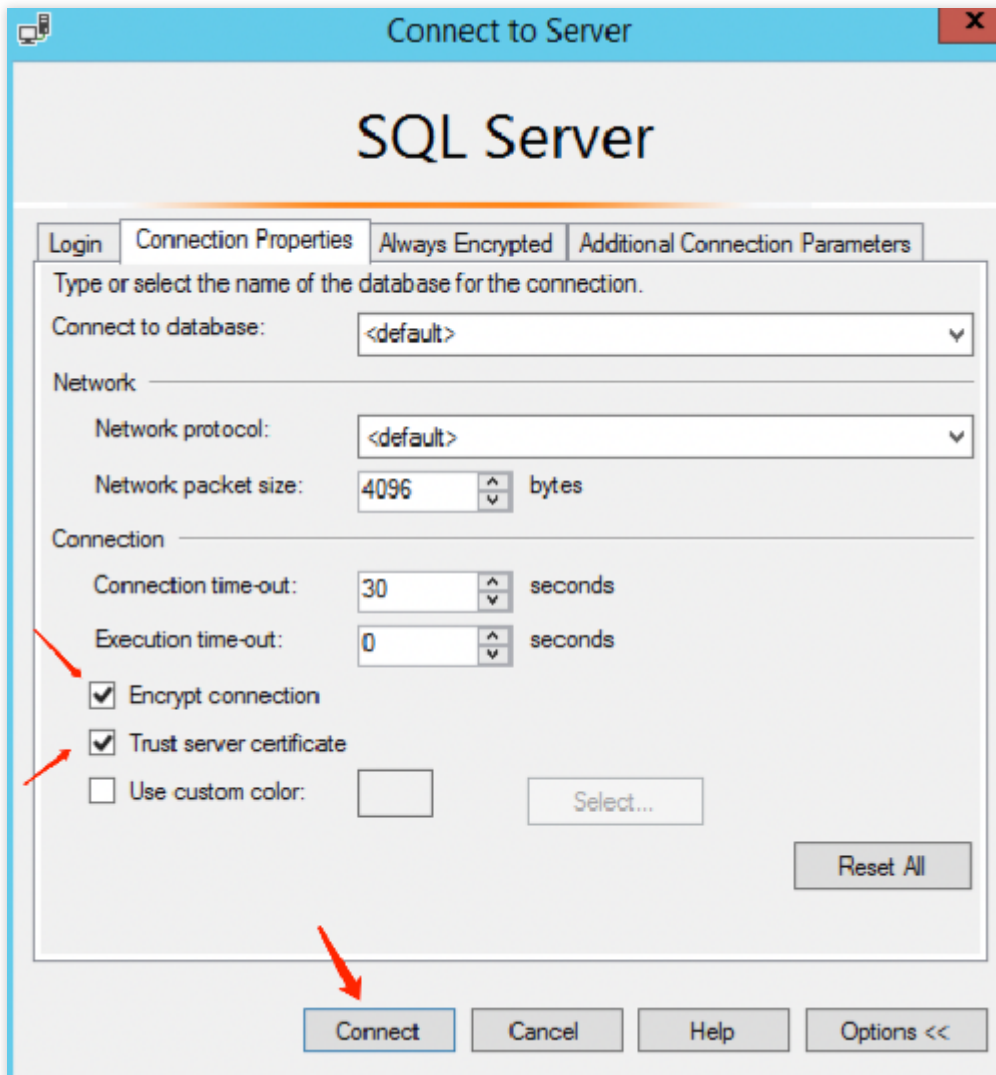
Scenario One: Encrypted Connection and Trusted Server Certificate

Step 1: Encrypted Connection Login

1. Open the SQL Server Management Studio client, and click **Option** in the lower right corner of the dialogue box.



2. On the **Connection Properties** tab, enable the **Encrypt Connection** and **Trust Server Certificate** options, then click **Connect**.

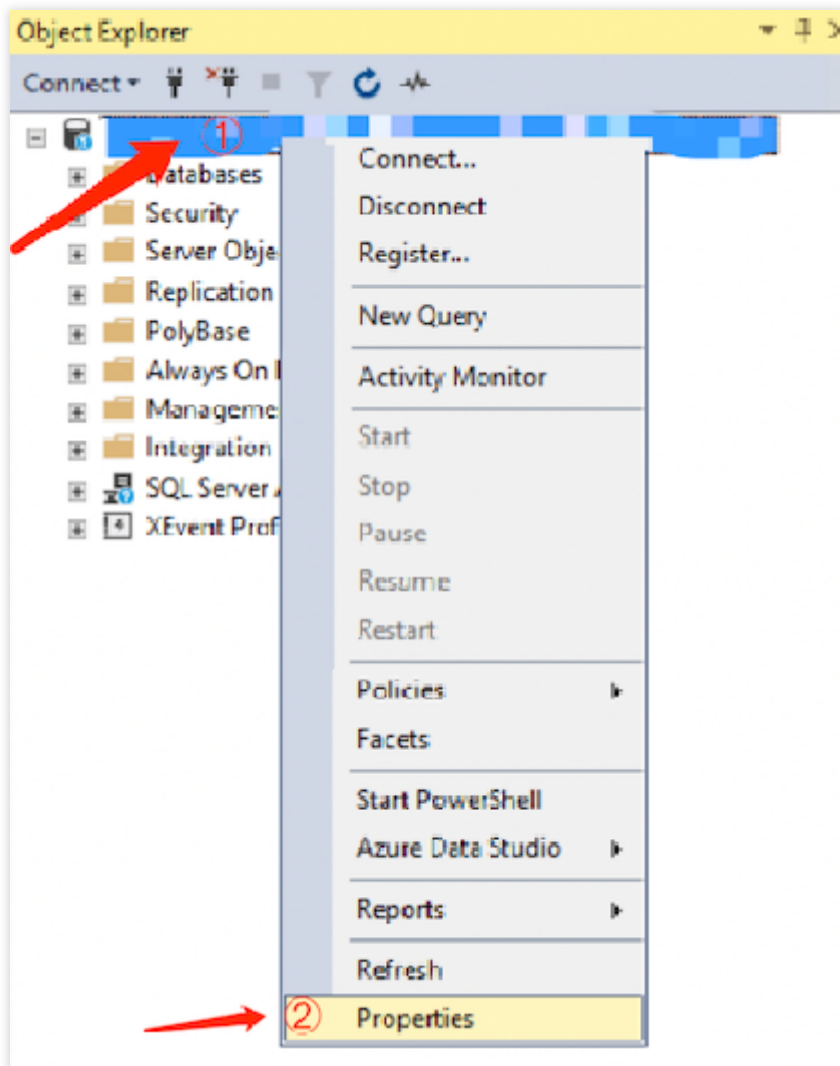


Step 2: Validate Whether the Connection is Encrypted

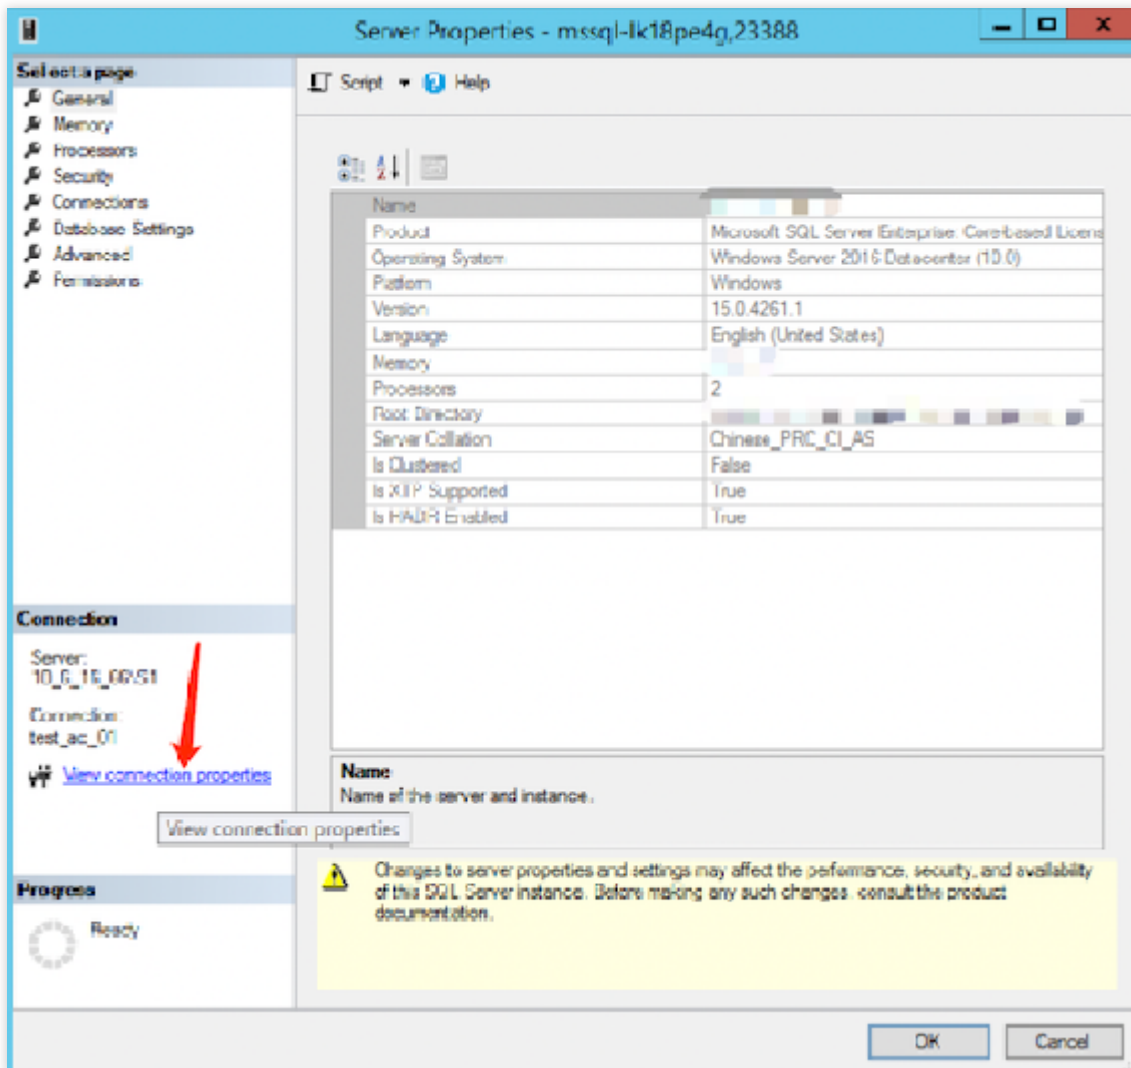
Method one: Via the SSMS client interface

Method two: Via commands

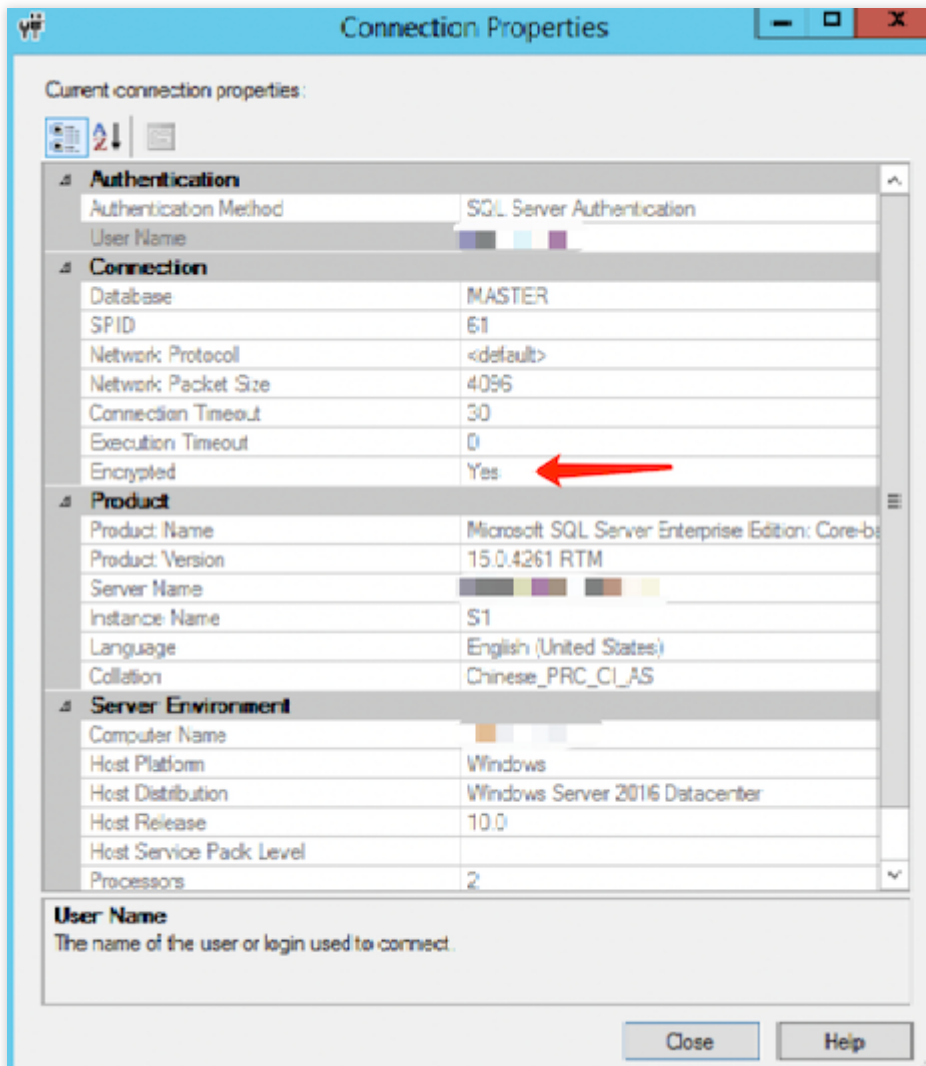
1. Once the connection is successful via SSMS, right-click on the instance and select **Properties**.




2. On the service properties page that pops up, click **View connection properties** on the left.



3. On the connection property page, you can view whether the connection has been encrypted.



Execute the following query command. If the query result is TRUE, it indicates that the connection is encrypted. If the result is FALSE, it signifies that the connection is not encrypted.



```
SELECT ENCRYPT_OPTION FROM SYS.DM_EXEC_CONNECTIONS WHERE SESSION_ID = @@SPID;
```

Scenario Two: Encrypted connection and untrusted server certificate

Step 1: Download the CA Certificate

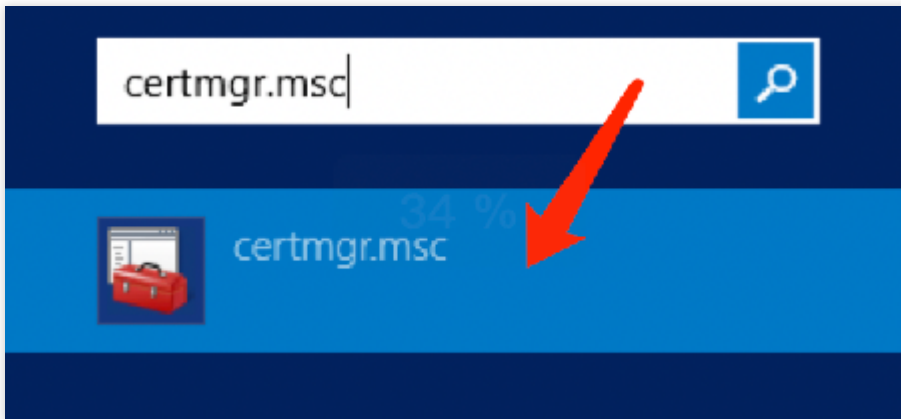
After the SSL encryption is enabled, on the target **instance details page** > **Data Security** > **SSL Encryption**, click [Download CA certificate](#).

Step 2: Clients import the CA certificate

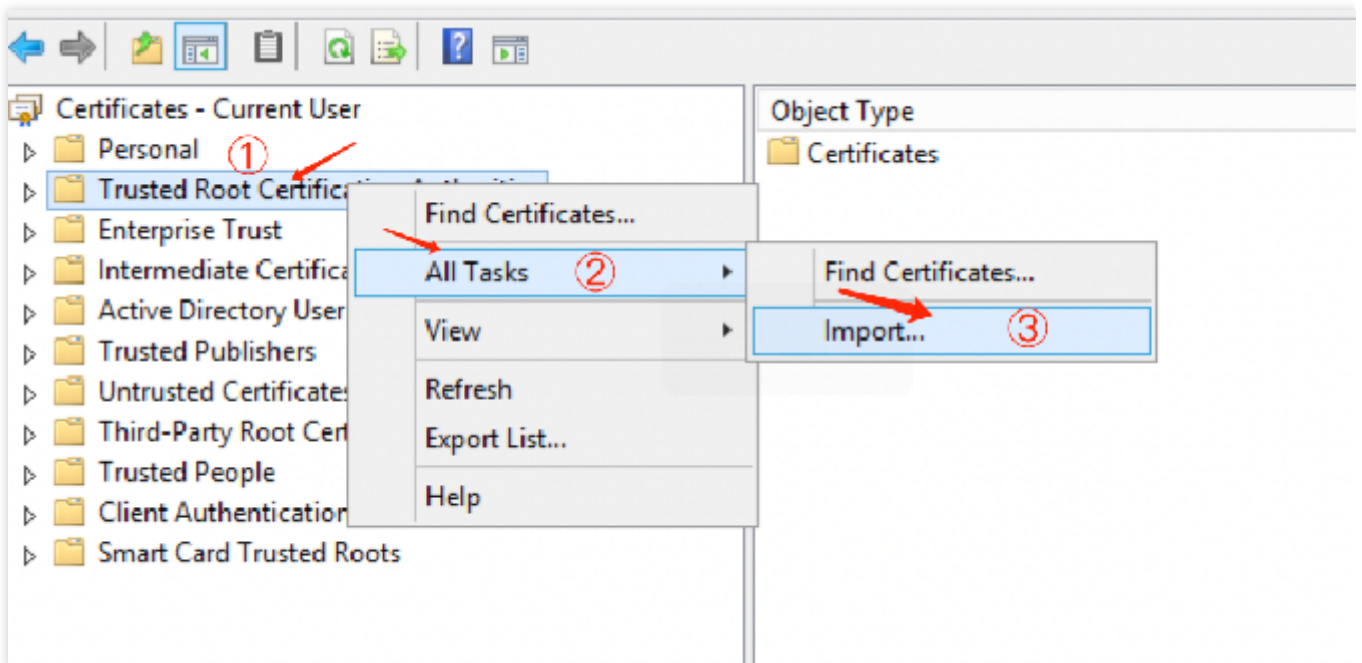
Method one: Through local client interface

Method two: Via command

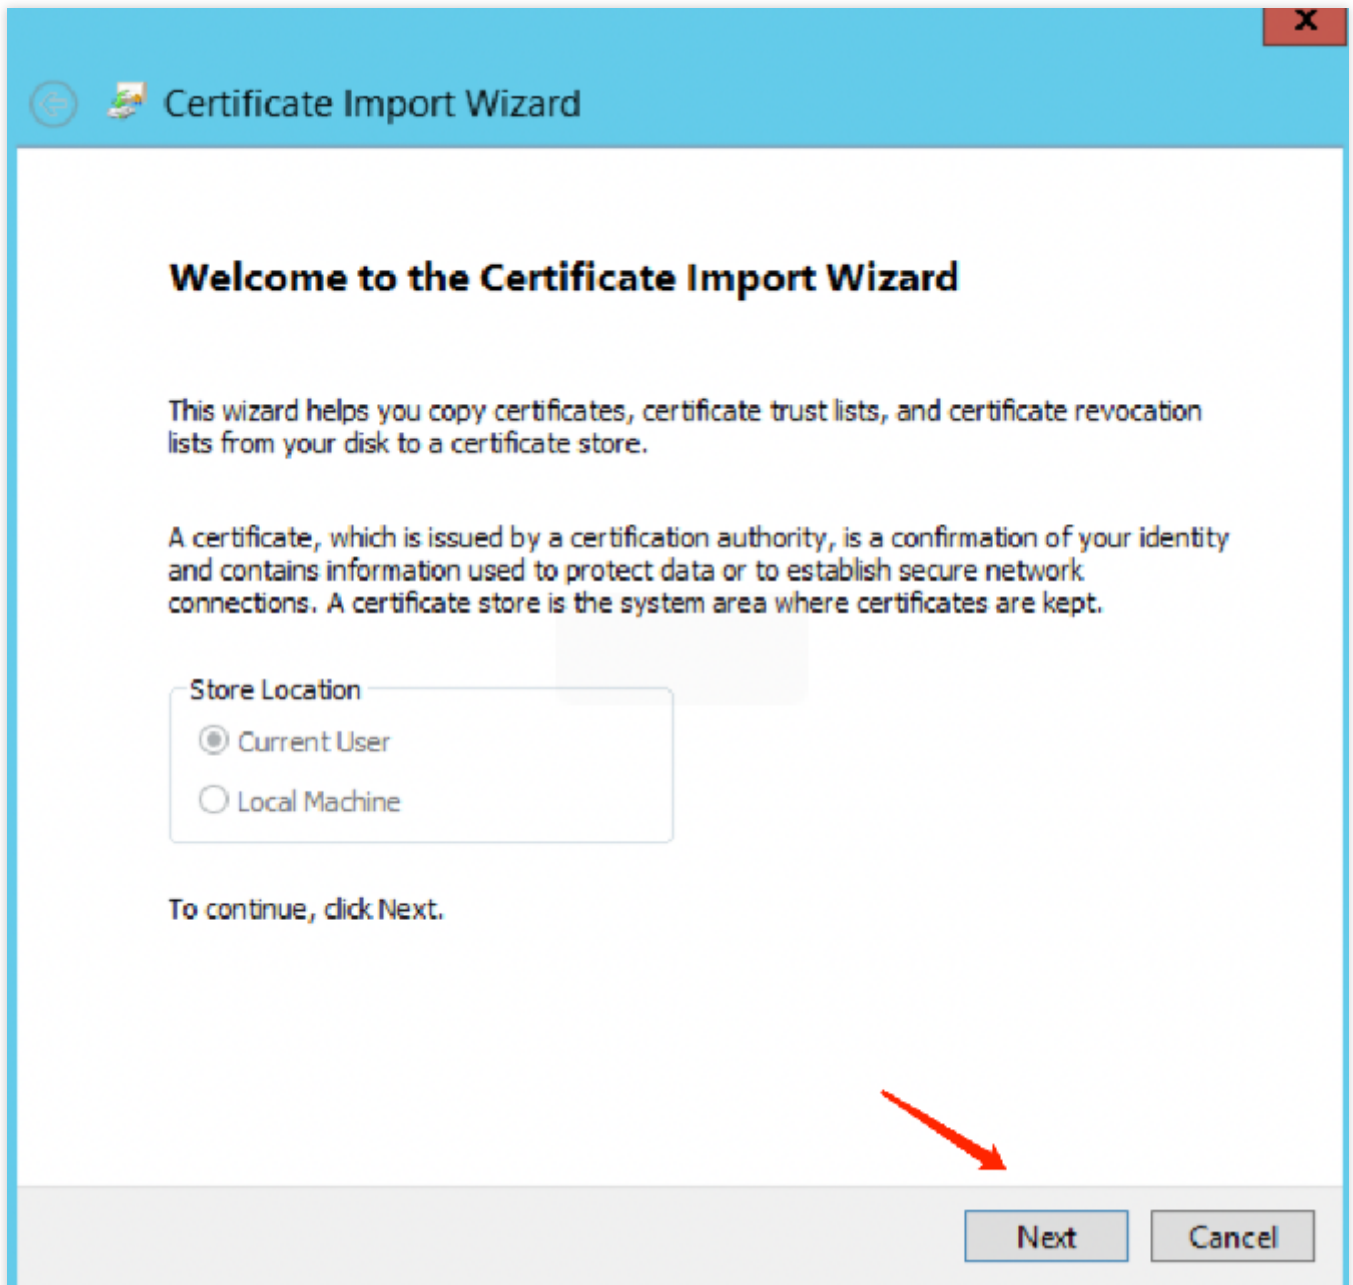
1. Click the search box in the lower-left corner of the desktop, and input certmgr.msc to open the Certificate Manager.



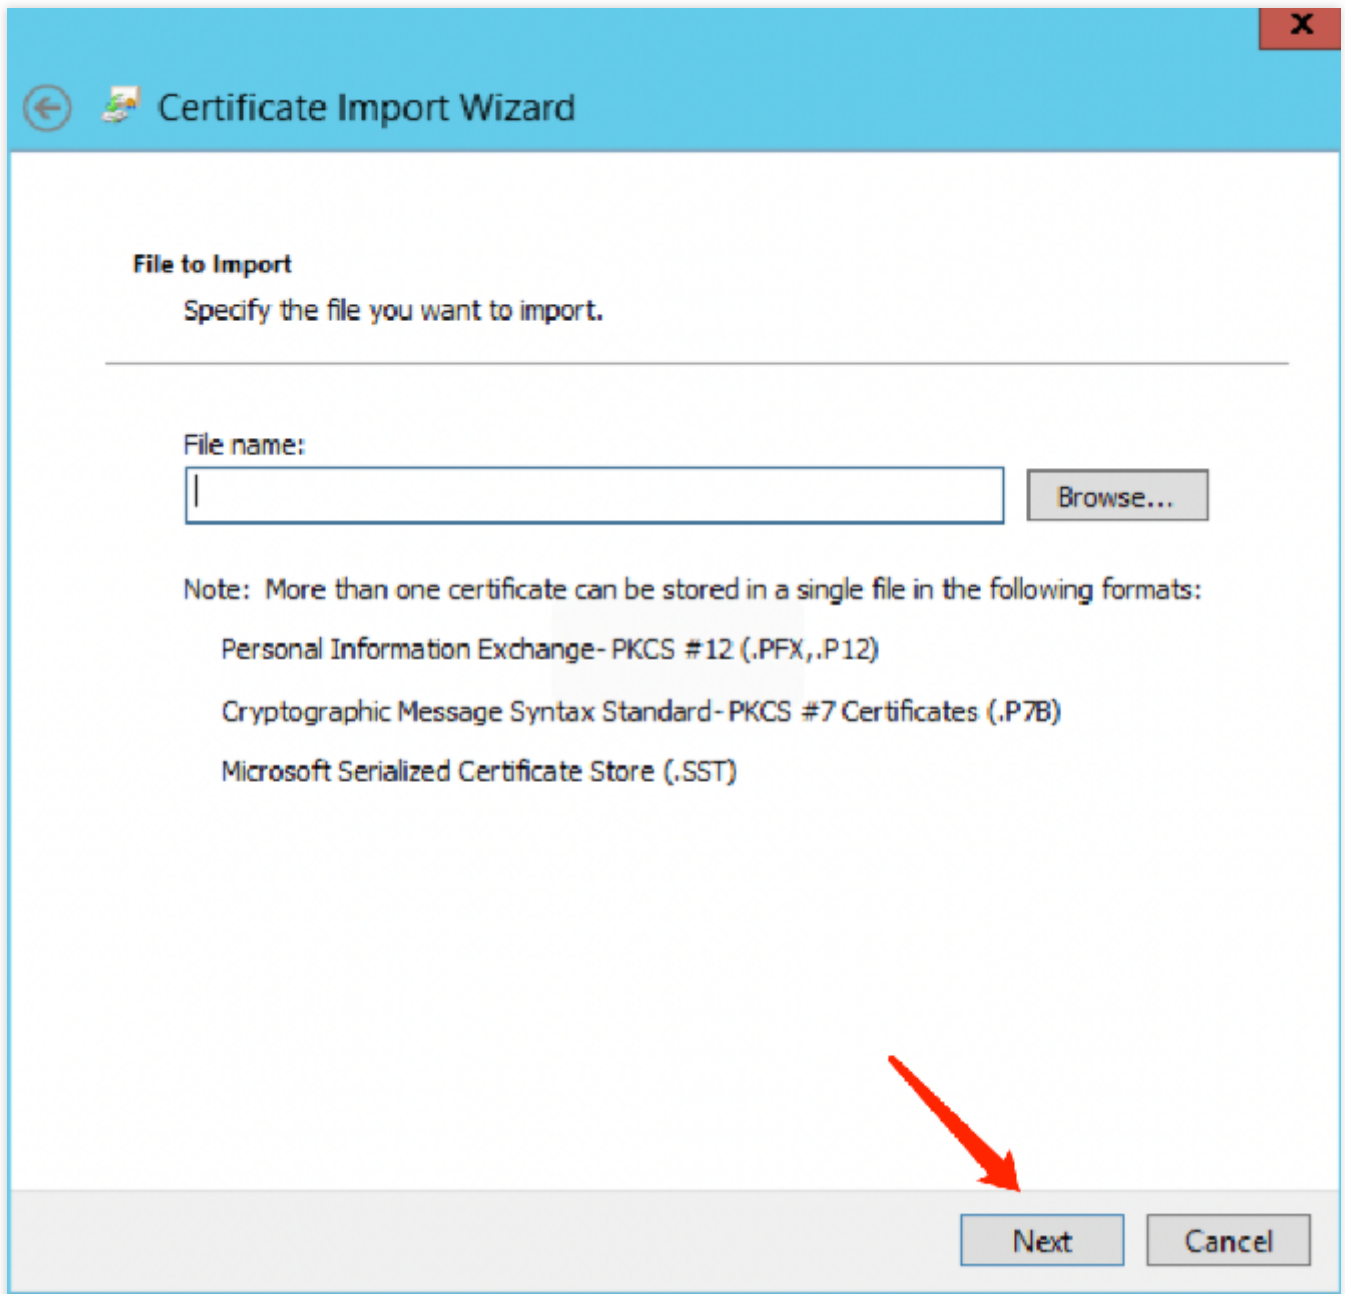
2. In the certmgr dialog box, right-click **Trusted Root Certification Authorities**, and choose **All Tasks > Import**.



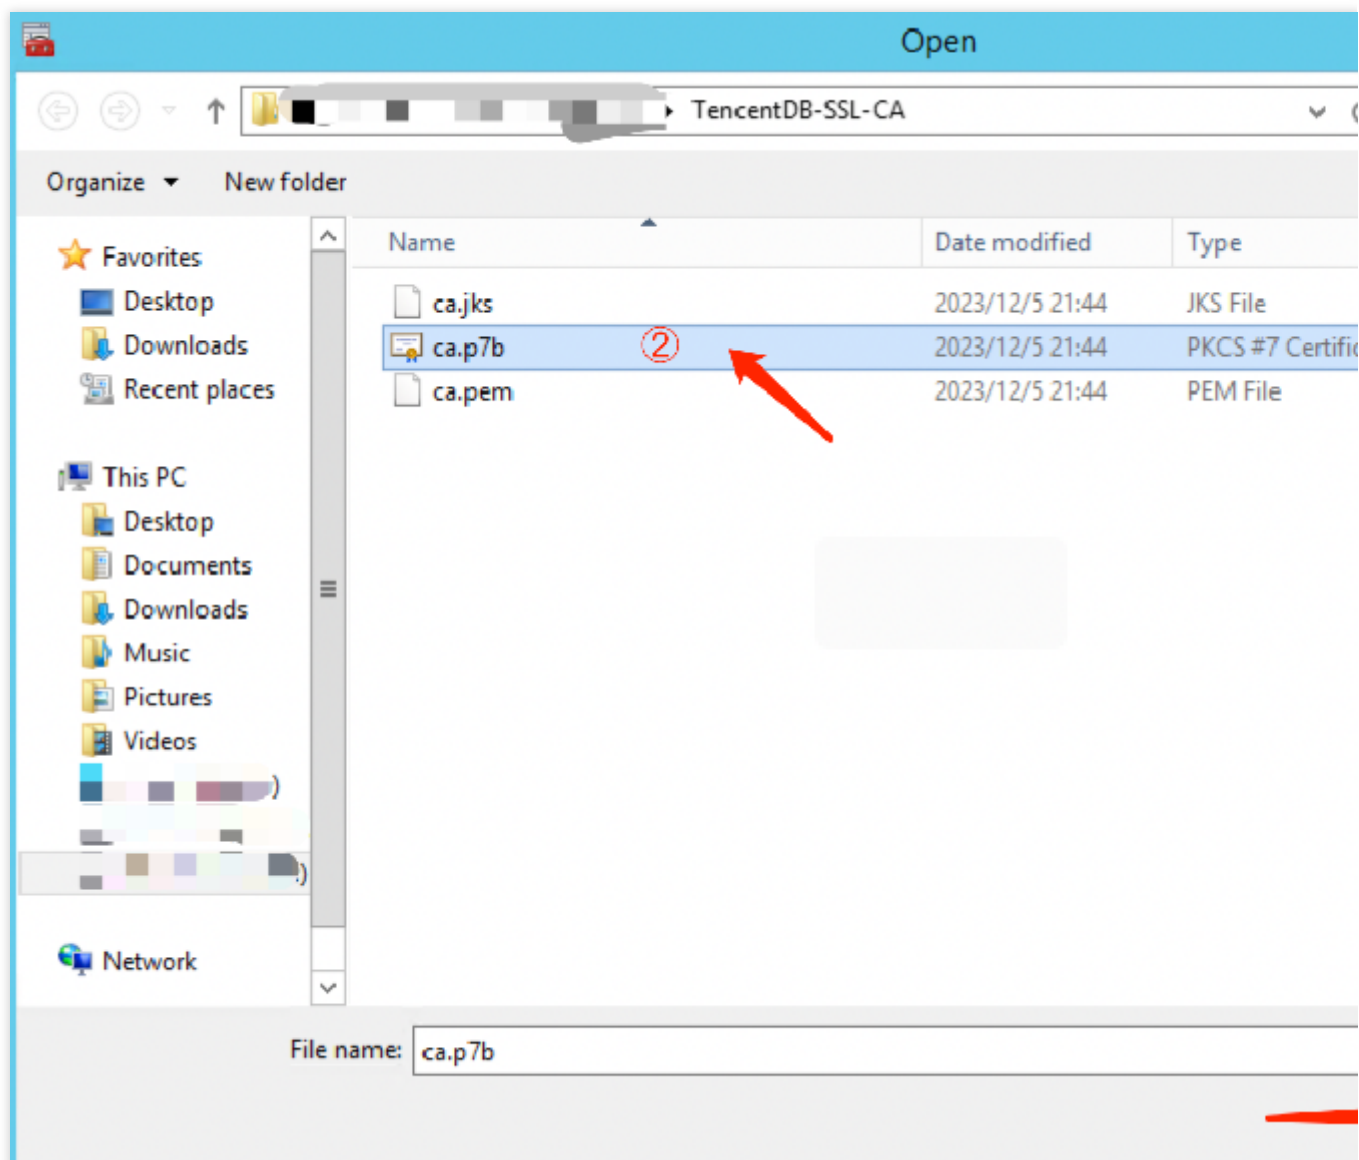
3. Click **Next**.



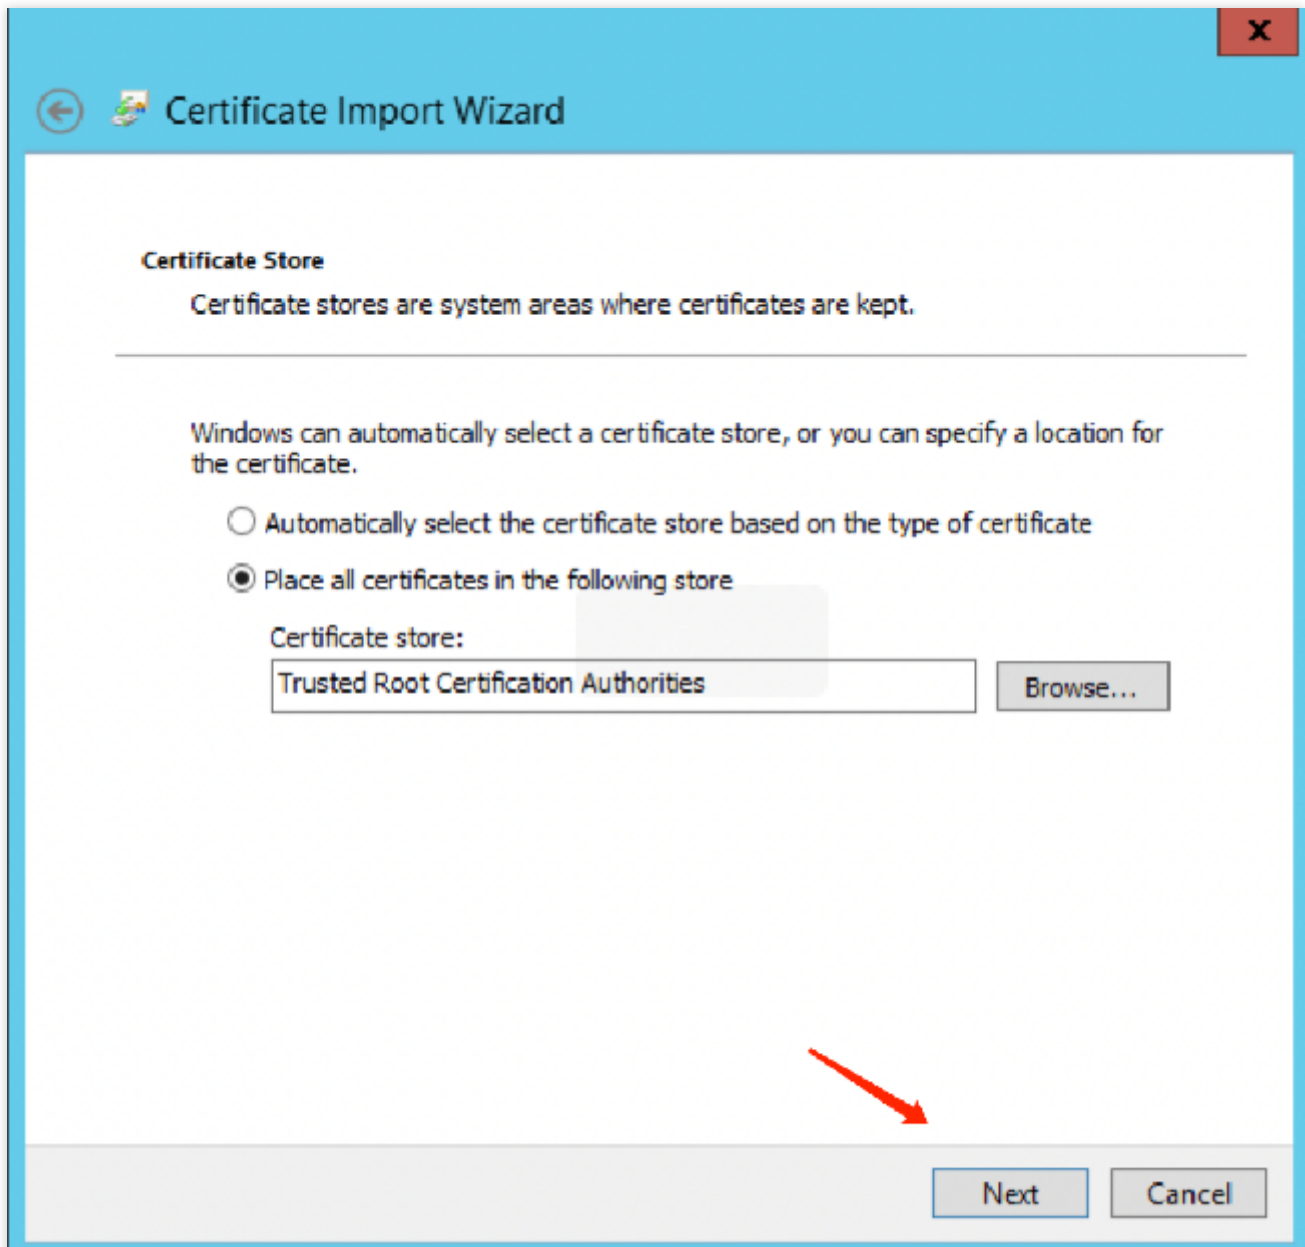
4. On the **Certificate Import Wizard** page, click **Browse** to import the downloaded SSL CA certificate, and then click **Next**.




5. Select the locally downloaded TencentDB-SSL-CA certificate file, and click **Open**.



6. Select the location to store the certificates according to personal needs, then click **Next**, followed by clicking **Done**.



You can also import the SSL CA certificate by commands. For instance, you can execute the following command to import the certificate via CMD or PowerShell.




```
CERTUTIL -addstore -enterprise -f -v root "ca.p7b"
```

Step 3: Log in through an encrypted connection

1. Open the hosts file located at C:\\Windows\\System32\\drivers\\etc.
2. Add IP mssql-***** at the end of the hosts file.

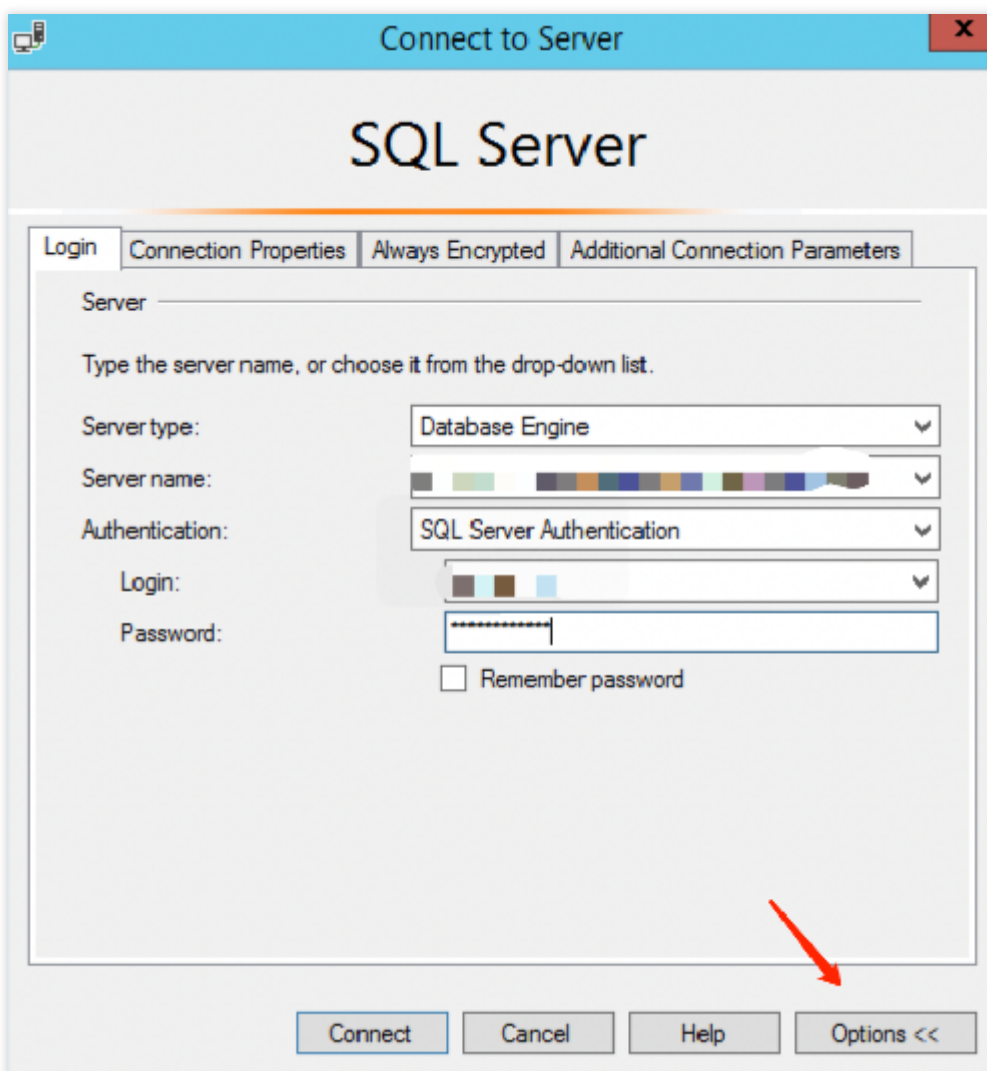
```
# localhost name resolution is handled within DNS itself.  
#          0.1          localhost  
#          ::1          localhost  
  
IP mssql-*****|
```

**Note:**

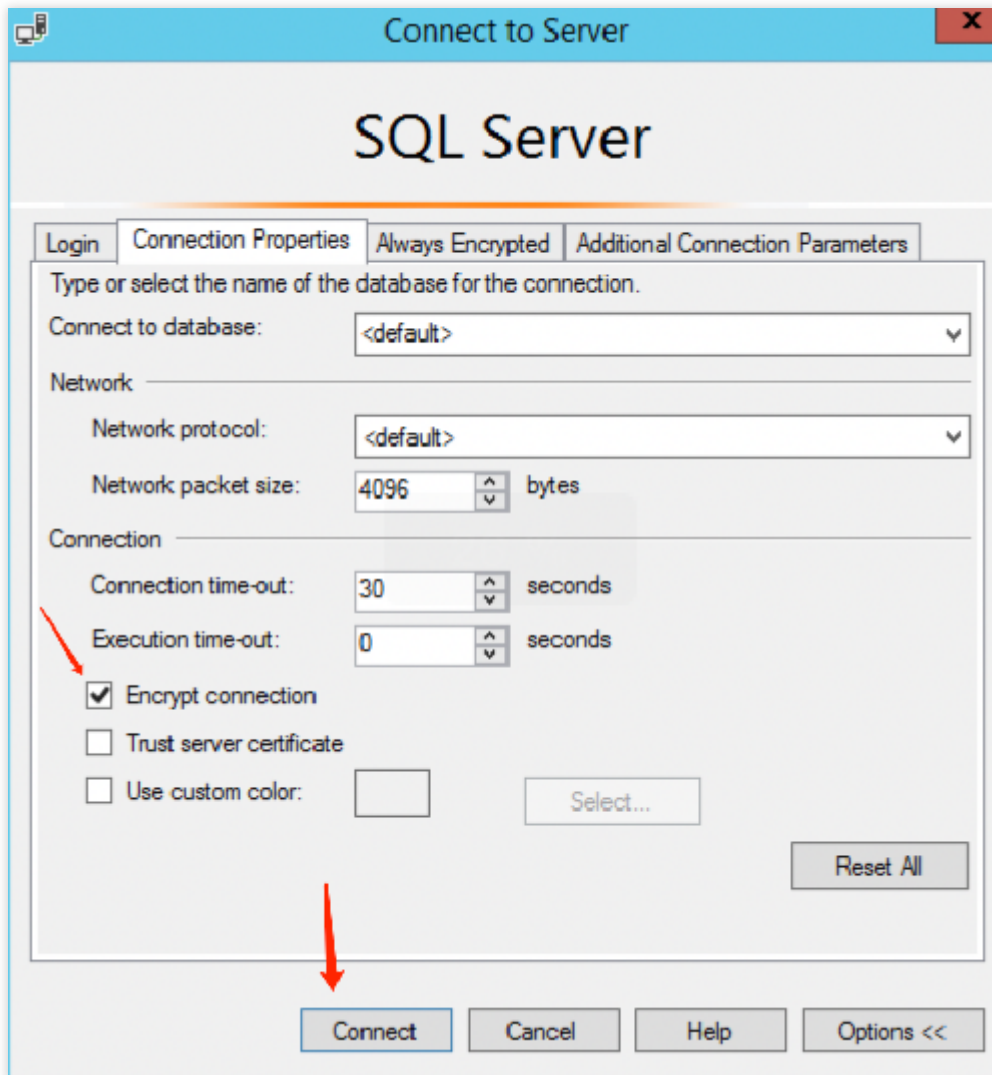
IP: Substitute it with the private or public IP address of the the corresponding instance.

mssql-*****: Replace it with the corresponding instance ID.

3. Open the SQL Server Management Studio client, click on the **Option** in the lower right corner of the dialogue box.



4. On the **Connection Properties** tab, check the **Encrypt Connection**, and then click **Connect**.

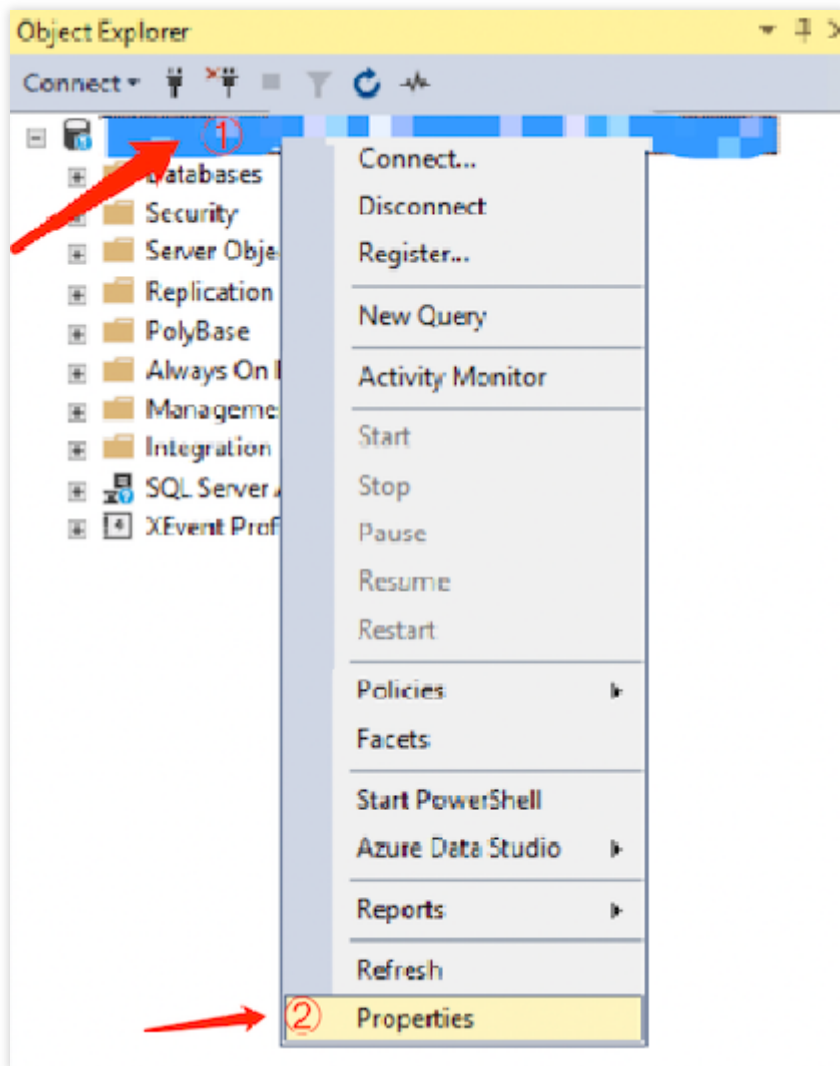


Step 4: Verify if the connection has been encrypted

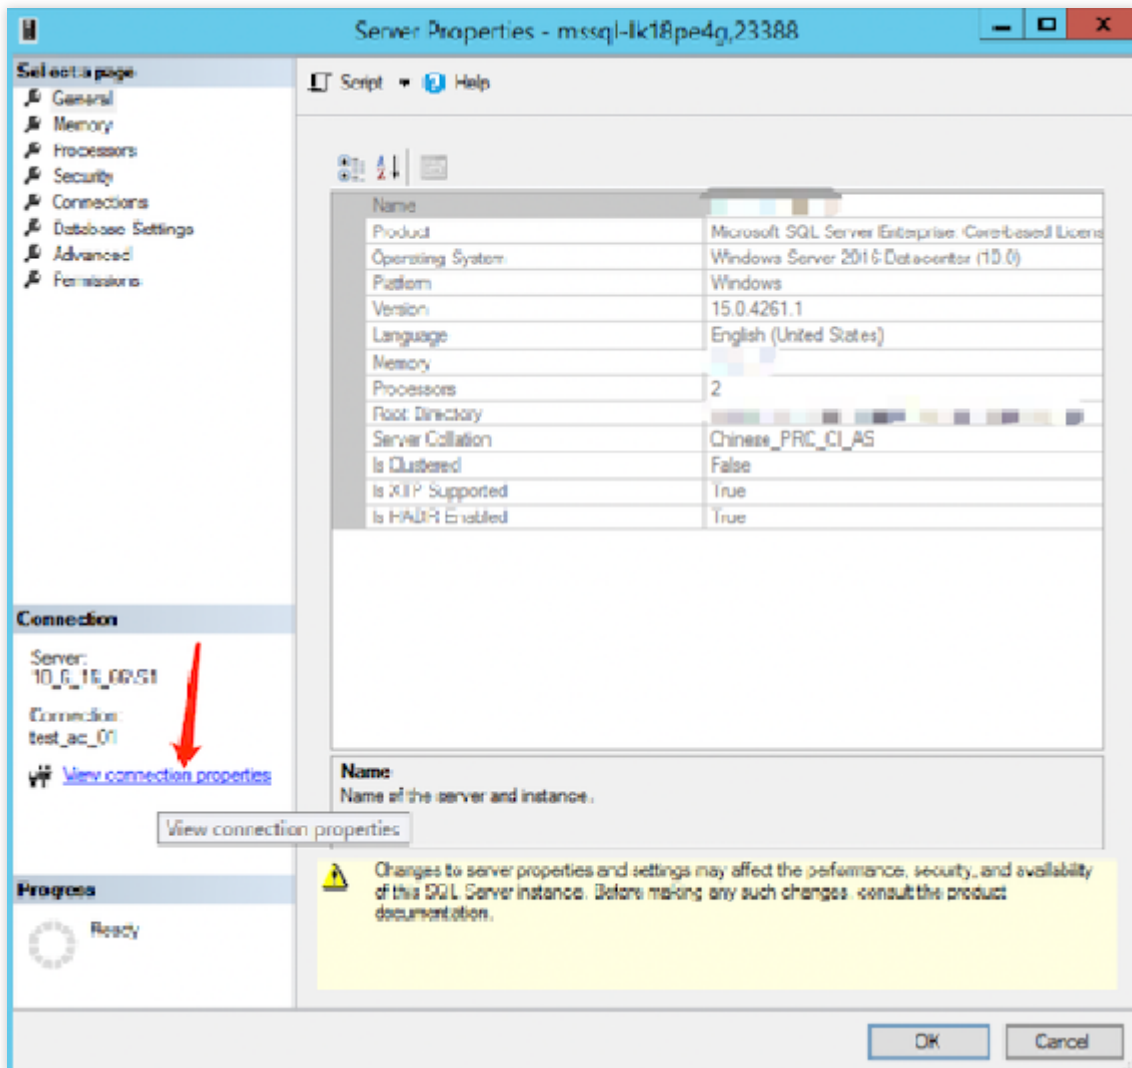
Method one: Via the SSMS client interface

Method two: Via commands

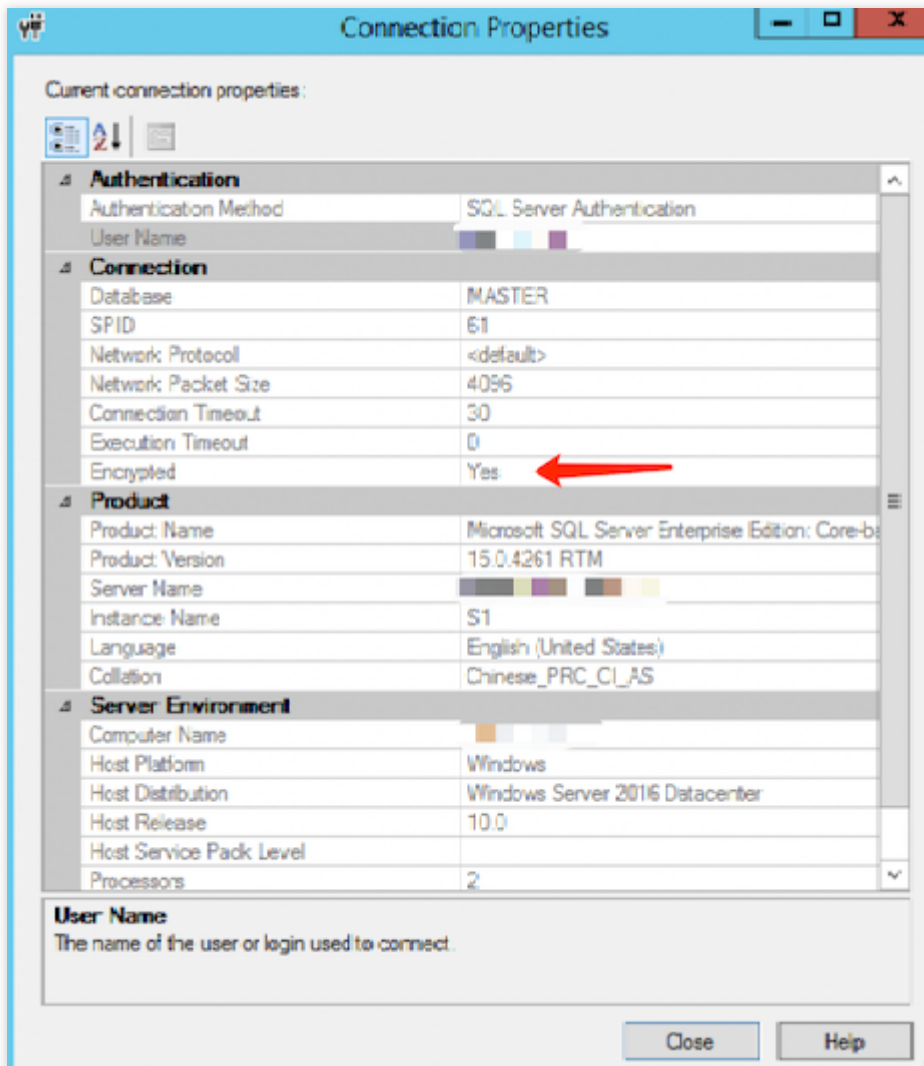
1. Once the connection is successful via SSMS, right-click on the instance and select **Properties**.




2. On the service properties page that pops up, click **View connection properties** on the left.



3. On the connection property page that pops up, you can view whether the connection has been encrypted.



Execute the following query command. If the query result is TRUE, it indicates that the connection is encrypted. If the result is FALSE, it signifies that the connection is not encrypted.



```
SELECT ENCRYPT_OPTION FROM SYS.DM_EXEC_CONNECTIONS WHERE SESSION_ID = @@SPID
```

Updating the validity period of the certificate

Note:

The SSL certificate has a validity period (one year). You need to manually update the certificate before it expires to ensure that client programs using SSL encrypted connections can continue to connect normally.

During the update of SSL Certificates, your database instance will be restarted to load the new SSL Certificates. Please ensure that your business has a reconnection mechanism.

1. Log in to the [TencentDB for SQL Server console](#).
2. Select **Region**. Click **Instance ID** of the target SSL or **Manage** in the **Operation** column.
3. Choose **Data Safety** > **SSL Encryption** on the **Instance Management** page, then click **Update Certificate** after the validity of SSL Certificates.
4. Select the execution time in the pop-up window, and click **Confirm**.

Update SSL Encryption Certificate

Execution Time ⓘ

Immediate execution	During maintena
---------------------	------------------------

Maintenance Time: 00:00-06:00 (modify on the "I page)

☒ Updating SSL certificate will restart your insta certificate. Please ensure that your business t reconnection mechanism.

OK

Cancel

Execution time:

Immediate execution: SSL Certificates will promptly be updated upon you click **OK**.

During maintenance time: Update SSL Certificates within the instance maintenance time. [Modification of the instance maintenance time](#) can be done on the instance details page.


Disabling SSL encryption

Note:

During the process of disabling SSL encryption, your database instance will be restarted. Please ensure that your business has a reconnection mechanism.

1. Log in to the [TencentDB for SQL Server console](#).

2. Select **Region**. In the instance list, click **Instance ID** of the instance that needs to disable SSL or **Manage** in the **Operation** column.
3. Select **Data Security** > **SSL Encryption** on the **Instance Management** page, then click the button to disable the function after **SSL Encryption Status**.



4. Select the execution time in the popup, and click **OK**.

Execution time:

Immediate execution: SSL encryption will be disabled immediately upon you click **OK**.

During maintenance time: Disable SSL encryption within the instance maintenance time. [Modification of the instance maintenance time](#) can be done on the instance details page.

Appendix 1: Changes for Instance-related Status during SSL Settings

During SSL settings, if the chosen execution time falls within the **maintenance time**, the system will check at a 10-minute interval whether the SSL settings related operations coincide with the maintenance time. If the operations coincide with the maintenance time, the system commences the deployment or termination of SSL, and promptly reboots the service upon completion.

Note:

If a user schedules an SSL operation during the maintenance time, and the designated instance is in a non-operational status prior to reaching this window, further tasks cannot be initiated from this instance until the SSL encryption operation has concluded.

Example

Assume that the maintenance time is set from 17:00 to 18:00. If a user configures to enable SSL encryption within this maintenance time at 17:05, the system initiates the asynchronous scheduling task at 17:10. Once successful deployment of SSL Certificates is accomplished, the service restarts immediately.

Explanation for Instance Status Changes

The status change of instances chosen to be executed during the maintenance time is as follows:

Prior to 17:05, the instance status was **in operation**;

Between 17:05 and 17:10, the status of the primary instance and the RO Replica instance is **in the process of SSL operation**;

After 17:10, the status of the primary instance and the RO replica instance is **in the process of task execution**; Not until the asynchronous task is completed do the state of both the primary instance and RO replica instance revert back to **Running**.

The status changes of instances for which the execution time is set to immediate execution are as follows:

After an SSL operation task is initiated via the console, the status of both the primary instance and the RO replica instance changes to **Task Execution In Progress**. It only reverts back to **Running** after the asynchronous task is completed.

Explanation of the SSL Operation Status Change of Instances

The changes in the instance operation status when SSL encryption is initiated are as follows:

The execution time is set to immediate execution: **Not activated**-> **Activating**-> **Activated**.

Execution time is within the maintenance time: **Not activated**-> **Waiting for activation within the maintenance time window**-> **Activating**-> **Activated**.

The changes to the instance operation status when the certificate is updated are as follows:

The execution time is set to immediate execution: **Activated** -> **updating** -> **initiated** .

The execution time is within the maintenance time: **Not initiated**-> **Waiting for updates within the maintenance time window**-> **Updating**-> **Activated**.

The changes in the instance operation status are as follows when SSL is disabled:


Execution time is set to immediate execution: **Enabled**-> **Being disabled**-> **Disabled**.

The execution time is within the maintenance time: **Enabled**-> **Waiting for disabling in maintenance time window**-> **Being disabled**-> **Disabled**.

FAQs

How can you view whether the SSL encryption has been enabled for the current connection?

You can query the sys.dm_exec_connections dynamic management view by using the following command, which will inform you whether the current connection has SSL encryption:



```
SELECT session_id,encrypt_option
FROM sys.dm_exec_connections;
GO
```

If the sys.dm_exec_connections dynamic management view query returns the session ID of the current connection and the encrypt_option value is true, then it demonstrates that the connection has successfully enabled SSL encryption.

Why was the client software able to establish a connection last year, but fails to establish a connection this year?

The SSL certificate has an expiration date, typically one year. The issue might be due to certificate expiration, which needs manual renewal. For detailed steps, refer to [Renewing the Certificate Validity Period](#).

Parameter Configuration

Setting Instance Parameters

Last updated : 2024-01-18 17:23:30

You can view and modify certain parameters and query parameter modification logs in the TencentDB for SQL Server console.

Notes

To ensure instance stability, only some parameters can be modified in the console. These parameters are displayed on the **Parameter Settings** page.

If the modified parameter requires instance restart to take effect, the system will ask you if you wish to restart. We recommend you do so during off-peak hours and ensure that your application has a reconnection mechanism.

Batch Modifying Parameters

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance management page.
2. On the instance management page, select the **Parameter Configuration > Parameter Settings** tab and click **Batch Modify Parameters**.

Parameter Settings

Modification Log

Batch Modify Parameters


Parameter Name	Instance Restart	Default Value	Curr
fill factor(%) ⓘ	No	0	80
max worker threads ⓘ	No	0	0
cost threshold for parallelism ⓘ	No	5	5

3. Select the target parameters and modify their values in the **Current Value** column. After confirming that everything is correct, click **OK**.

Parameter Settings Modification Log

Batch Modify Parameters

Parameter Name	Instance Restart	Default Value	Current Value
fill factor(%) ⓘ	No	0	80
max worker threads ⓘ	No	0	0
cost threshold for parallelism ⓘ	No	5	5
max degree of parallelism ⓘ	No	0	0
optimize for ad hoc workloads ⓘ	No	0	0 ▼
min server memory(MB) ⓘ	No	0	1024
blocked process threshold(s) ⓘ	No	0	0

 Modify the current values, with 0 values modified

4. In the pop-up window, select the **Execution Mode** and click **OK**.

Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately.

If you select **During maintenance time**, the parameter modification task will be executed and take effect during the [instance maintenance time](#).

Modify Parameters ✕

A total of 2 parameters will be modified. Continue?

Parameter Name	Original Value	New Value
max degree of parallelism	0	12
optimize for ad hoc workl...	0	1

Execution Mode Immediate execution During maintenance time

OK Cancel

Modifying One Parameter


- Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID to access the instance management page.
- On the instance management page, select the **Parameter Configuration > Parameter Settings** tab, select the row of the target parameter, and click



in the **Current Value** column to modify its value.

Parameter Settings Modification Log

Batch Modify Parameters

Parameter Name	Instance Restart	Default Value	Current Value
fill factor(%) ⓘ	No	0	80  Click to modify

- Enter the target parameter value as prompted in the **Current Value** column and click



to save the change. You can click



to cancel the operation.

Parameter Name	Instance Restart	Default Value	Current Value
fill factor(%) ⓘ	No	0	80

4. In the pop-up window, select the **Execution Mode** and click **OK**.

Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately.

If you select **During maintenance time**, the parameter modification task will be executed and take effect during the [instance maintenance time](#).

Modify Parameters

A total of 1 parameters will be modified. Continue?

Parameter Name	Original Value	New Value
fill factor(%)	80	81

Execution Mode

Immediate execution

During maintenance time

OK

Cancel


Viewing Parameter Modification Log

Last updated : 2024-01-18 17:23:30

You can view parameter modification logs in the TencentDB for SQL Server console.

Directions

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID to access the instance management page.
2. View the parameter modification logs in **Parameter Configuration > Modification Log**.



Monitoring and Alarms

Viewing Monitoring Charts

Last updated : 2024-01-18 17:23:30

To make it easier for you to view and stay up to date with how instances work, TencentDB for SQL Server provides a wide variety of performance monitoring metrics and convenient monitoring features (custom view, time comparison, merged monitoring metrics, etc.).

This document describes how to view monitoring chart info in the console.

Note:

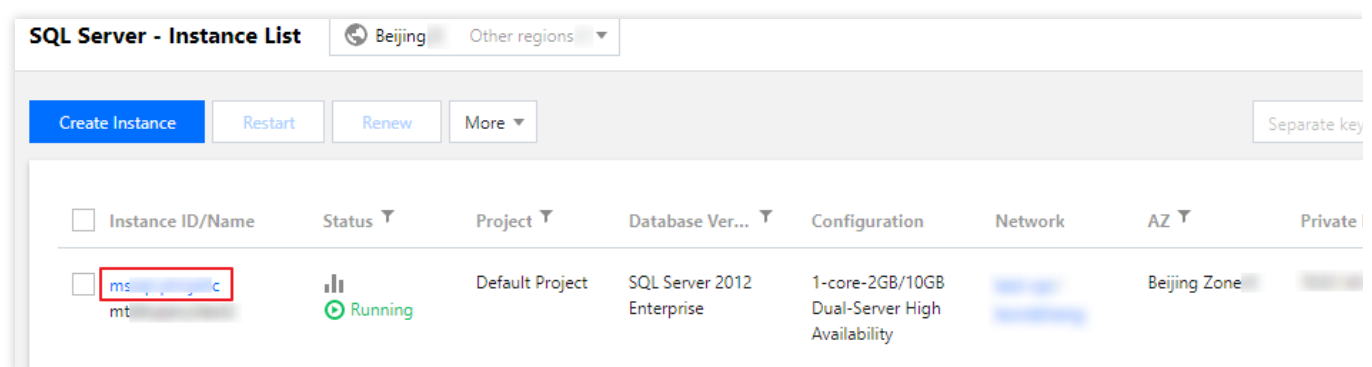
If the number of tables in a single instance exceeds one million, database monitoring may be affected. Make sure that the number of tables in a single instance is below one million.

Types of instances for monitoring

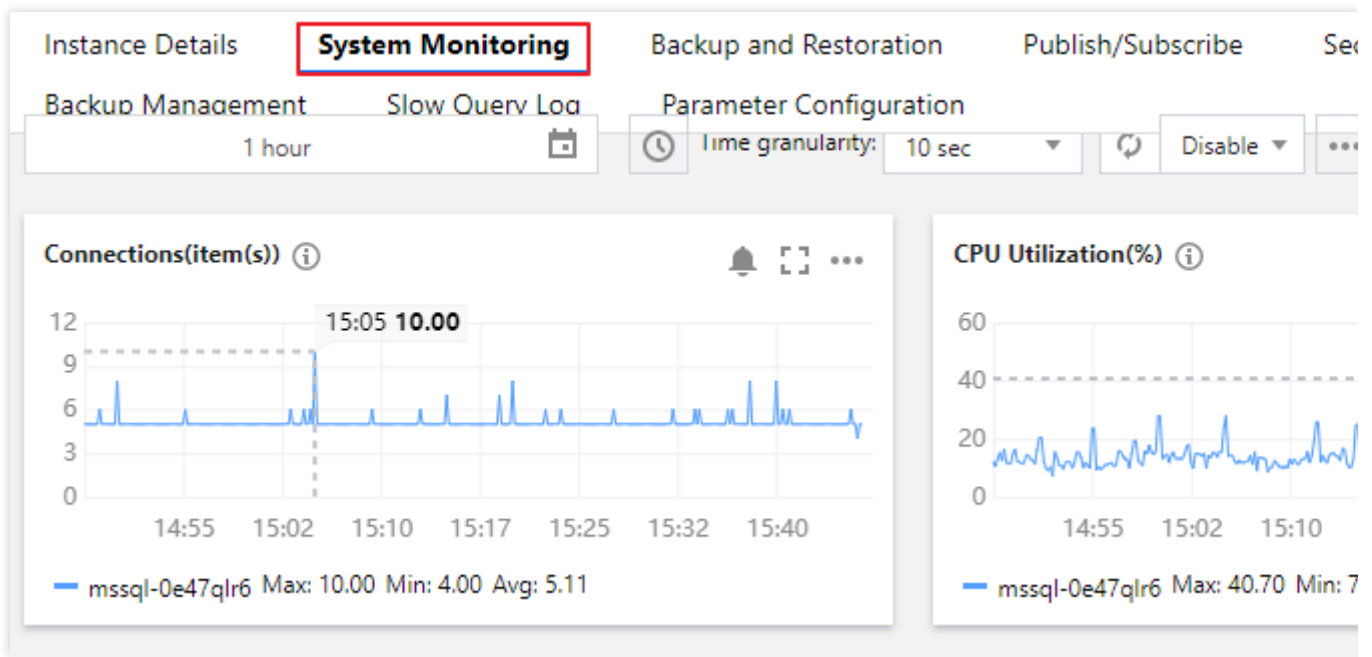
TencentDB for SQL Server primary and read-only instances can be monitored, and each instance is provided with a separate monitoring view for easy query.

Viewing monitoring data

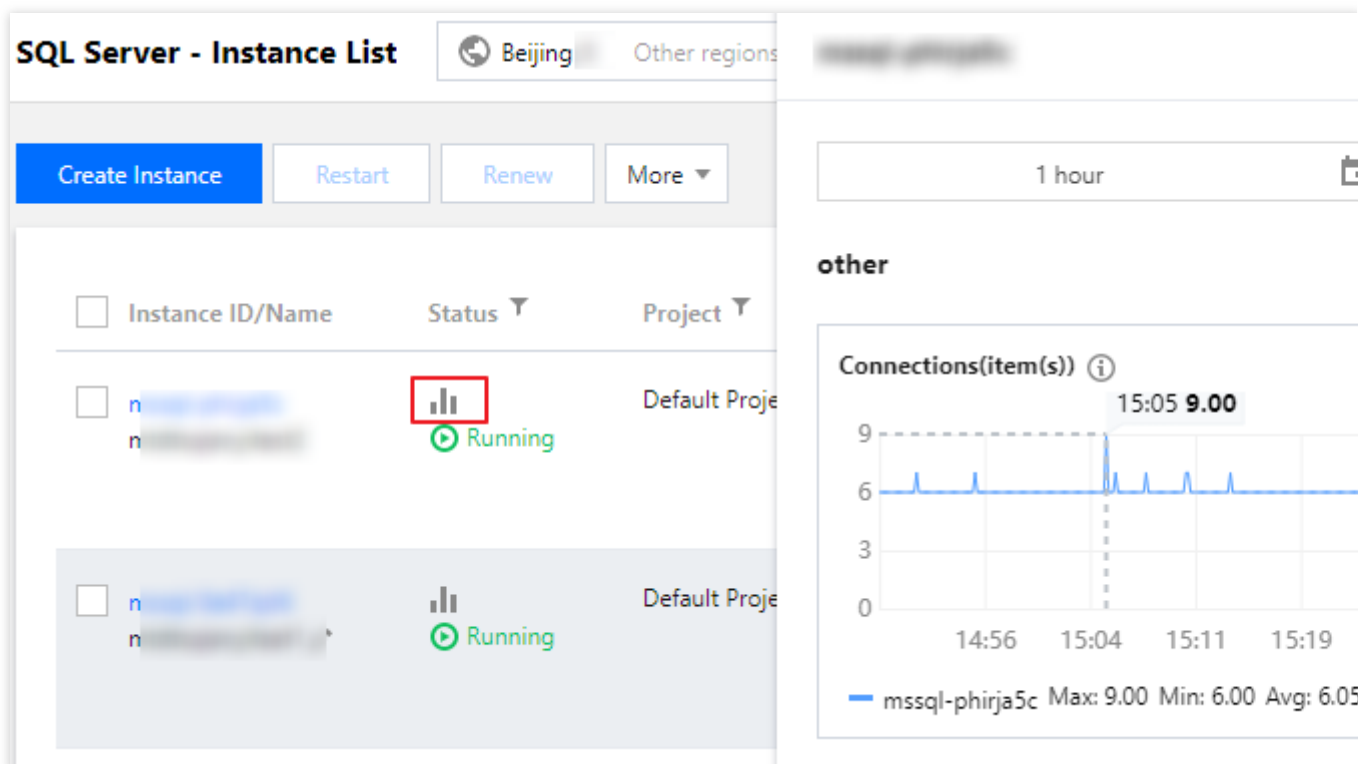
1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select the **System monitoring** page to view monitoring metrics.

**Note:**

On the instance list, you can also click the monitoring icon of the target instance to quickly view its monitoring status.



Displaying a chart in full screen


You can display a single metric in full screen for a clearer preview of metric data.

1. On the [System Monitoring](#) page, you can click icon



on the right of the corresponding metric to display the metric in full screen.

2. After the preview data is displayed in full screen, you can click the X in the upper right corner to close the full screen display window.



Exporting data

You can export the desired metric data individually.

On the [System Monitoring](#) page, you can click icon




on the right of the corresponding metric to export data or pictures of the metric to the local system.



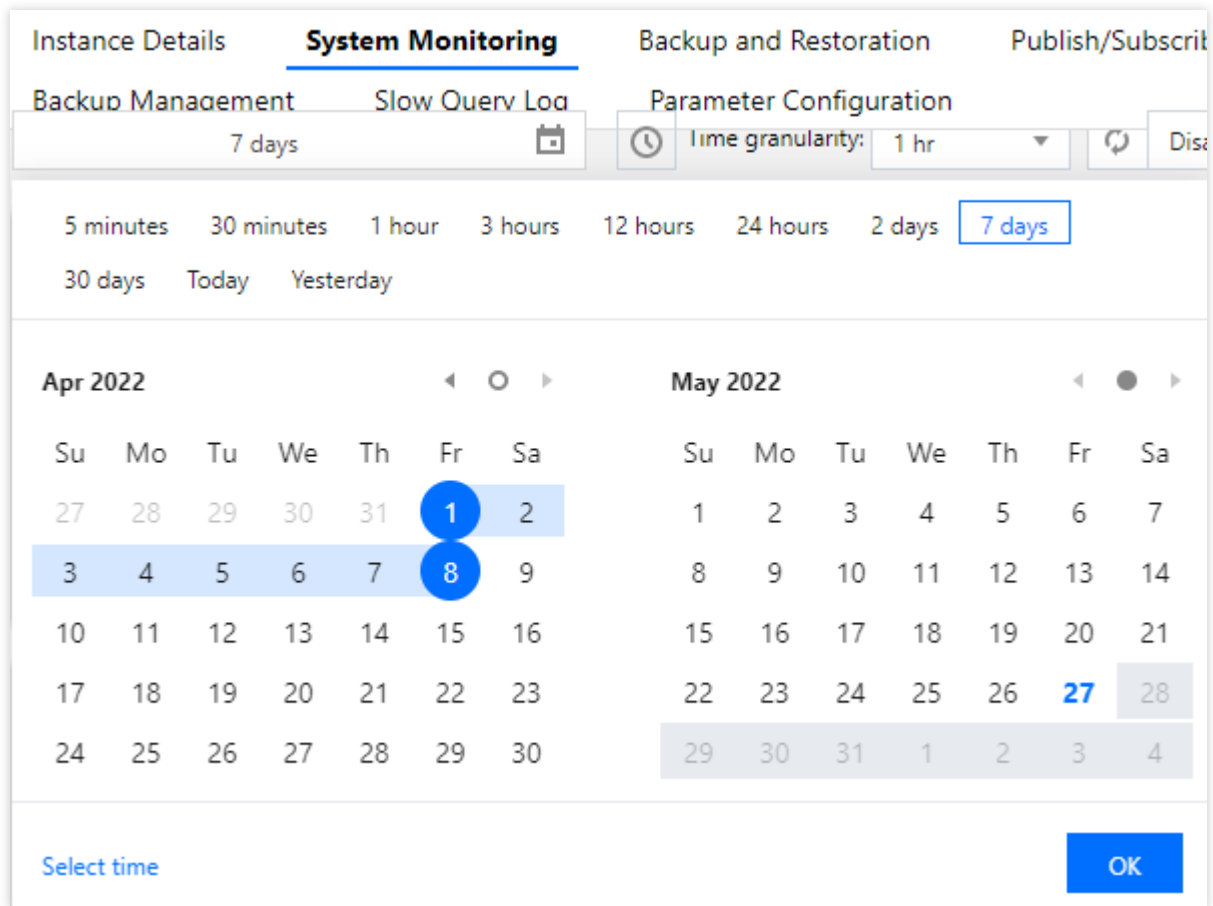
Selecting monitoring time range

You can select or customize a time range to query the monitoring over this time period.

1. On [System Monitoring] page, click the time box.




2. In the pop-up window, you can select 5 minutes, 30 minutes, 1 hour, 3 hours, 12 hours, 24 hours, 2 days, 7 days, 30 days, today, yesterday, the start and end date, or the time range of them, then click **OK**.



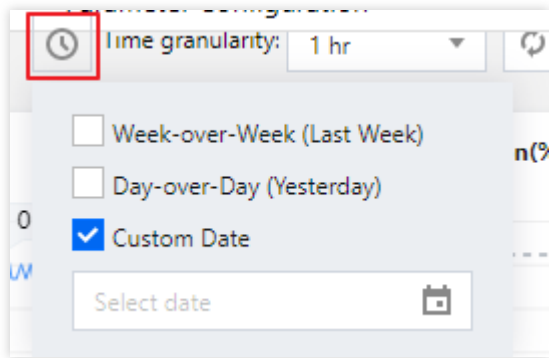
Adding time comparison

You can compare monitoring data from multiple time ranges by adding time comparisons.

1. On [System Monitoring] page, click **Add Time** icon behind time box.




2. Select **Week-over-Week (Last Week)**, **Day-over-Day (Yesterday)**, or **Custom Date** from the drop-down list, and click **Ok**.



Monitoring granularity

You can view instance monitoring at different time granularities within the selected time period.

On [System Monitoring] page, select the time period, and then you can select the desired time granularity in the drop-down list after **Time Granularity** to view the monitoring data.



Time Periods and Corresponding Chart Granularities

Time period	Time granularity
5 minutes	10 seconds, 1 minute
30 minutes, 1 hour	10 seconds, 1 minute, 5 minutes
3 hours	10 seconds, 1 minute, 5 minutes, 1 hour
12 hours, today, yesterday	1 minute, 5 minutes, 1 hour
24 hours, 2 days	1 minute, 5 minutes, 1 hour, 24 hours
7 days, 30 days	1 hour, 24 hours

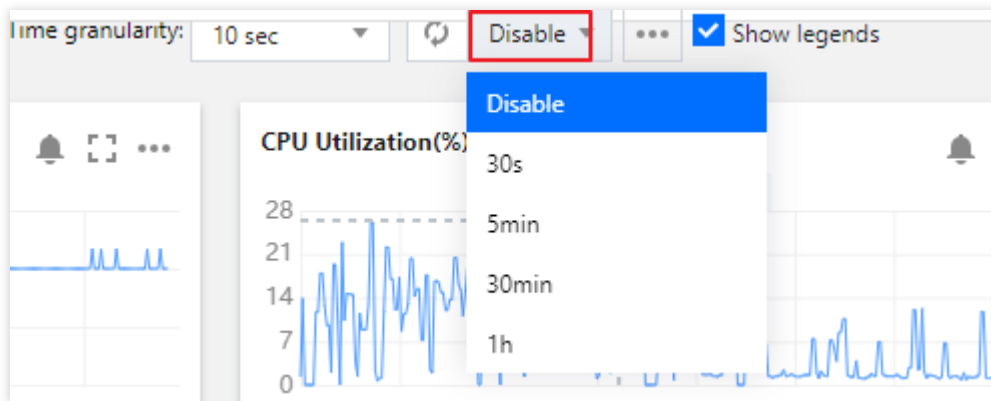
Setting refresh time

You can set the refresh time on the system monitoring page (disabled by default) to observe changes of instance monitoring in real time.

On the [System Monitoring](#) page, click the drop-down button after



to set time frequency of data refresh (30 seconds, 5 minutes, 30 minutes, and 1 hour are supported).



Monitoring Metrics

Last updated : 2024-07-30 16:09:09

This document describes the common monitoring metrics supported by TencentDB for SQL Server.

Monitoring Metrics

TencentDB for SQL Server supports monitoring 43 common metrics of SQL Server. You can collect statistics of other metrics by configuring the performance counters of SQL Server Management Studio (SSMS).

Note:

To view monitoring metrics at the system level, [contact us](#).

Classification	Metric Name	Unit	Descriptions
CPU	CPU Usage	%	Percentage of actual CPU consumption Note : When the user's business QPS is less than 5 or there is no business database, the CPU will be displayed as 0. For instances with cloud disk architecture, the CPU metrics are collected from the CPU of the machine where the instance is deployed. In cases of smaller specifications (such as 2C or 4C), there may be some fluctuations.
Memory	Memory Usage	MB	Total memory usage
	Max Memory	KB	Max memory
	Memory Usage Rate	%	Memory usage rate
	Internal System Locked Memory	KBytes	Internal system locked memory This monitoring metric can be viewed only in 2014 and later versions.
	Internal System Consumed Memory	KBytes	Internal system consumed memory This monitoring metric can be viewed only in 2014 and later versions.
Storage	Disk Throughput	KBytes/s	The total size of data written and read to the disk per second (for monitoring machines deployed

		with dual node cloud disk instances)
Disk Read Traffic	KBytes/s	Amount of data read from disk to memory per second (for monitoring machines deployed with dual node cloud disk instances)
Disk Write Traffic	KBytes/s	Amount of data written from memory to disk per second (for monitoring machines deployed with dual node cloud disk instances)
Disk IOPS	Times/sec	Number of disk read and write times
Number of Disk Read Times	Times/sec	Number of disk read times per second
Number of Disk Write Times	Times/sec	Number of disk write times per second
Used Storage	GB	Sum of storage space consumed by instance database files and log files
Remaining Storage Space	%	Percentage of remaining hard disk capacity, calculated as: (Purchased disk - Used storage space) / Purchased disk
Read IO Average Response Time	ms	The average time required to read data from the disk, measured in seconds, for monitoring the machines deployed with two-node cloud disk instances and system-level monitoring of machines deployed with two-node local disk instances.
Write IO Average Response Time	ms	The average time required to write data to the disk, measured in seconds, for monitoring the machines deployed with two-node cloud disk instances and system-level monitoring of machines deployed with two-node local disk instances.
IO Request Average Response Time	ms	Average response time per IO request (for system-level monitoring of machines deployed with dual node local disk instances)
Number of Read IOs Per Second on Disk	Times/sec	Number of read IOs per second on disk (for system-level monitoring of machines deployed with dual node local disk instances)
Number of Write IOs	Times/sec	Number of write IOs per second on disk (for

	Per Second on Disk		system-level monitoring of machines deployed with dual-node local disk instances)
	Total Disk IOPS	Times/sec	Total IOPS on disk (for system-level monitoring of machines deployed with dual-node local disk instances)
	Disk Queue Length	Count	Disk queue length (for system-level monitoring of machines deployed with dual-node local disk instances)
Network	Inbound Traffic	KB/s	Sum of inbound packet sizes for all connections
	Outbound Traffic	KB/s	Sum of outbound packet sizes for all connections
	Average Network I/O Delay	ms	Average network I/O delay time
	RO Sync Delay Time	Sec	Data sync delay time between primary instances and read-only instances
Connection	Number of Connections	Connections	Average number of databases connected by users per second
	Number of Logins Per Second	Times/sec	Number of logins per second
	Number of Logouts Per Second	Times/sec	Number of logouts per second
Access	Slow Queries	Queries	Number of queries running longer than one second
	Number of Requests	Times/sec	Average number of requests per second
	Number of SQL Compilations	Times/sec	Average number of SQL compilations per second
	Number of SQL Recompilations	Times/sec	Average number of SQL recompilations per second
	Number of Full-Table Scans for SQL Per Second	Times/sec	Number of full-table scans for SQL per second
	Number of Transactions	Times/sec	Average number of transactions per second

	Number of Blockings	Blockings	Number of current blockings
	Execution Cache Hit Rate	%	Each SQL statement has an execution plan, the hit rate of the execution plan.
	Buffer Cache Hit Rate	%	Data cache (memory) hit rate
Lock	Number of Lock Requests	Times/sec	Average number of lock requests per second
	Number of Latch Waits Per Second	Times/sec	Number of latch waits per second
	Average Latency on a Lock Wait	Ms	Average wait time of each lock request resulting in wait
	Number of Deadlocks	Number/sec	Number of lock requests causing deadlocks per second
	Blocking Report Generation Threshold	Minute	The specified threshold for generating blocking reports. If the threshold is exceeded, a blocked process report will be generated. By default, blocking process reports are not generated.
Others	Number of User Errors	Times/sec	Average number of errors per second

Setting Alarm Policies

Last updated : 2024-01-18 17:23:30

Scenario

You can create alarm policies to trigger alarms and send alarm notifications when the Tencent Cloud service status changes. The created alarm policies can determine whether an alarm needs to be triggered according to the difference between the monitoring metric value and the given threshold at intervals.

You can take appropriate precautionary or remedial measures in a timely manner when the alarm is triggered by changed product status. Therefore, properly created alarm policies can help you improve the robustness and reliability of your applications. For more information on alarms, see [Creating Alarm Policy](#) in Cloud Monitor.

To send an alarm for a specific status of a product, you need to create an alarm policy at first. An alarm policy is composed of three compulsory components, that is, the name, type and alarm triggering conditions. Each alarm policy is a set of alarm triggering conditions with the logical relationship "OR", that is, as long as one of the conditions is met, an alarm will be triggered. The alarm will be sent to all users associated with the alarm policy. Upon receiving the alarm, the user can view the alarm and take appropriate actions in time.

Note:

Make sure that you have set the default alarm recipient; otherwise, no notifications will be sent based on the default alarm policy of TencentDB.

Directions

1. Set alarm policies


1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top of the page, and click the instance ID in the instance list or **Manage** in the **Operation** column of the target instance to enter the instance management page.

SQL Server - Instance List

Beijing Other regions

Create Instance Restart Renew More

Separate key

<input type="checkbox"/>	Instance ID/Name	Status	Project	Database Ver...	Configuration	Network	AZ	Private
<input type="checkbox"/>	ms-...-c mt	 Running	Default Project	SQL Server 2012 Enterprise	1-core-2GB/10GB Dual-Server High Availability		Beijing Zone	

3. On the instance management page, select **System Monitoring**, and click any monitoring metric to enter alarm policy page.

4. On the **Create Alarm Policy** page, set the policy name, policy type, project, alarm object, and trigger condition, etc. After confirming that everything is correct, click **Complete**.

Policy Name: It can contain up to 30 characters.

Remarks: It can contain up to 100 characters.

Monitoring Type: Select **Cloud Product Monitoring**.

Policy Type: Select TencentDB for SQL Server.

Project: It is used for alarm policy classification and limit management, and it doesn't have strong binding to SQL Server instance project.

Alarm Object: The instance to be associated with the policy alarm. You can find the desired instance by selecting the region where it is located or searching for its ID.

Trigger Condition: You can select **Select template** or **Configure manually** and **Use preset trigger condition**

(The common trigger conditions of the alarm policy for the corresponding cloud products are automatically set according to the preset template of the system). An alarm trigger is a semantic condition composed of metric, comparison, threshold, statistical period, and duration. For example, if the metric is disk utilization, the comparison is $>$, the threshold is 80%, the statistical period is 5 minutes, and the duration is two statistical periods, then the data on disk utilization of a database will be collected once every five minutes, and an alarm will be triggered if the disk utilization exceeds 80% for two consecutive times.

Configure Alarm Notification: You can select a preset or custom notification template. Each alarm policy can be bound to three notification templates at most. For more information, see [Setting Alarm Notification](#).

Basic Info

Policy Name

It can contain up to 30 characters

Remarks

It can contain up to 100 characters

Monitoring Type

Cloud Product Monitoring

Policy Type

CDB (SQL Server)

Project

DEFAULT PROJECT

2 exist. You can create 298 more static threshold policiesThe current account has 0 policies for dynamic alarm thresholds, and 20 more policies can be created

Alarm Policy

Alarm Object

Instance ID

Select object

Trigger Condition

☐ Select template

☒ Configure manually (☒ Apply preset trigger conditions ⓘ)

Metric Alarm

When meeting

any

 of the following metric conditions, the metric will trigger an alarm.

Threshold Type ⓘ

☒ Static

☐ Dynamic ⓘ

▶ If

CPUUtilization

(statistical period

>

80

%

at 1 consecutive

then

Alarm once a hour

ⓘ

🗑

2. Associate objects

After the alarm policy is created, you can associate alarm objects with it. When an alarm object satisfies an alarm trigger condition, an alarm notification will be sent.

1. In the [Alarm Policy](#) list, click the name of an alarm policy to enter the alarm policy management page.
2. Click **Add Object** in the **Alarm Object** section.

Alarm Object

Add Object

Disassociate

Disassociate All

Regions that have no instances bound to alarm policy are not displayed

<input type="checkbox"/> ID/Name	Status
----------------------------------	--------

3. In the pop-up dialog box, select the alarm objects to be associated with, and click **OK**.

Setting Alarm Notification

Last updated : 2024-01-18 17:23:30

This document describes how to set an alarm notification template in the console. When creating an alarm policy on the alarm policy configuration page, you can select an existing template or create a new one for easy configuration. You can also quickly apply the same template for multiple policies.

For more information on relevant limits and other descriptions, see [Creating Notification Template](#).

Directions

1. Configure an alarm notification template

1. Log in to the [CM console](#) and select **Notification Template** on the left sidebar.
2. Click **Create**, enter information on the **Create Notification Template** page, confirm that everything is correct, and click **Complete**.

Basic Info

Template Name *

It can contain up to 30 Chinese characters, letters, di

Notification Template ⓘ

☒ Alarm Trigger ☒ Alarm Recovery

Notification Language

English

Notifications (Fill in at least one item)

User Notification

You can add a user only for receiving messages.

Recipient Object

User

Add User

Delete

Notification Cycle

☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat ☒ Sun

Notification Period

00:00:00 ~ 23:59:59 ⓘ

Receiving Channel

☒ Email ☒ SMS

Add User Notification

API Callback ⓘ

API Callback URL

Enter a URL accessible over public networks as the API callback address (domain name or IP[:port]/path), e.g. https://example

Delete View Usage Guides ⓘ

Notification Cycle

☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat ☒ Sun

Notification Period

00:00:00 ~ 23:59:59 ⓘ

Add API Callback

It supports pushing to the WeCom group robotTry Now ⓘ

Ship to CLS

☐ Enable ⓘ

Please select region

Select a logset

Select a log topic

Create Log Topic ⓘ

Complete

2. Configure an alarm notification

1. Log in to the [CM console](#) and select **Alarm Policy** on the left sidebar.
- 2 Click **Create** and select up to three alarm notification templates in **Configure Alarm Notification** in the **Configure Alarm Policy** window.
2. Click **OK** after configuring the alarm policy and notification.

Configure Alarm Notification To add an alarm recipient (group), you need to select a notification template or create one below. You can click

Notification Template

Select template Create Template

1 selected. 2 more can be selected.

Notification Template Name	Included Operations
默认 Notification 🔗	Recipient: 1 Edit Recipient

Advanced Configuration (Optional)
Auto Scaling ☒ After enabling, the auto scaling policy can be triggered when reaching alarm condition.

Complete

3. After the configuration is completed, if an alarm is triggered by an exception, the system will send notifications to the recipients via the specified channels (email, SMS, and phone call).

More

[Setting Alarm Policies](#)

Viewing Alarm Records

Last updated : 2024-01-18 17:23:30

You can create alarm policies to trigger alarms and send notifications in the console when the TencentDB for SQL Server instance status changes. The created alarm policies can determine whether an alarm needs to be triggered based on the difference between the monitoring metric value and the given threshold at intervals.

You can view specific information of alarm records in the console, and quickly locate the problem for further troubleshooting by alarm messages.

This document describes how to view alarm records in the console.

For more information on the feature, see [Viewing Alarm Records](#).


Directions

1. Log in to the [CM console](#), and select **Alarm List** on the left sidebar.
2. You can view alarm records in the alarm list.


Note:

You can click icon to customize the displayed fields of alarm records in the list, including occurrence time, monitoring type, policy type, alarm object, alarm content, duration, alarm status, policy name, end time, alarm type, alarm reception, alarm channel, instance group, project, and network.

Select a time or a custom time range to view the target alarm records.







Click **Advanced Filter** to view the target alarm records.



Enter alarm object in search box to view the target alarm records.

Advanced Filter



Backup and Restoration

Backup

Backup Overview

Last updated : 2024-07-30 15:37:58

The backup feature can help you avoid accidental data loss in case of system hardware or instance failure. To protect your assets, TencentDB for SQL Server provides the data backup and restoration feature for you to archive data and restore data to local databases.

This document describes the backup feature.

Backup purposes

You can restore data from backups after a database or table is deleted maliciously or by mistake. This helps guarantee data security and avoid data loss and corruption.

Backup billing

Backups include data backups and log backups. For more information on billing, see [Backup Space Billing](#).

Automatic backup

TencentDB for SQL Server supports automatic backup, including non-archive backup and archive backup. Archive backup is a more flexible backup policy on the basis of non-archive automatic backup and does not need to retain additional new backups. You can set the automatic backup retention policy based on your business needs to easily manage the backup retention period and cycle. You can also flexibly set the number of retained backups for the specified cycle to implement long-term backup storage, or shorten the non-archive backup retention period to reduce storage costs.

Non-Archive Backup: The data backup retention period is customizable between 7 and 1,830 days. The log backup retention period is the same as the data backup retention period by default. You can set the automatic backup cycle. We recommend you back up your data at least twice a week. For more information, see [Setting Non-Archive Backup Retention](#).

Archive Backup: The archive backup retention period is 365 days by default and customizable between 90 and 3,650 days. It can only be longer than the non-archive backup retention period. The archive backup retention

frequency and start time can be customized. For more information, see [Setting Archive Backup Retention](#).

Manual backup

Instance backup and multi-database backup are supported. You can manually create a backup file at any time for any created database. The larger the backup file, the slower the manual backup. Generally, it takes about 5–120 minutes to complete a manual backup. For more information, see [Creating Manual Backup](#).

Cross-region backup

You can store backup files in another region, which enhances the regulatory compliance and disaster recovery capabilities and improves the data reliability. After this feature is enabled, cross-region backup will be triggered after the local default automatic backup is completed, that is, the default automatic backup will be dumped to the cross-region backup storage device. The cross-region backup retention period is customizable between 7 and 1,830 days. For more information, see [Cross-Region Backup](#).

Data backup

Utilizing the manual backup feature, you have the flexibility to either back up individual or multiple databases within an instance or opt for a comprehensive backup of all databases across the entire instance. Conversely, with the automatic backup feature, you can back up the entire instance. The retention period for manually backed-up files can be aligned with the retention policy of automatic backups or can be set to be aligned with the lifecycle of the instance. By default, the retention period for automatically backed-up files is set to 7 days, with the option to customize the retention period ranging from 3 to 1830 days, after which the backup files are automatically deleted. It is recommended to promptly download any backup files you wish to retain locally.

Log backup

The system automatically generates log backups (log files) every 10 minutes, which are then uploaded to cloud storage. These log files are available for download. The retention period for log backups is consistent with that of data backups, ranging from 3 to 1830 days. Upon reaching the expiration date, the backup set is automatically deleted.

Backup policy

Backup policies include instance backup and multi-database backup. The former backs up all databases in an instance, while the latter backs up selected databases.

Backup task settings

You can configure the global variables of manual and scheduled backups through **Backup Task Settings**. You can set **Upload Backup File** to **Archive file** or **Unarchived files** and select the primary or replica instance for backup. For more information, see [Setting Backup Task](#).

Backup file format

The **Upload Backup File** option allows you to specify the format of backup files (**Archive file** or **Unarchived files**). By default, backup files (i.e., .bak files) will be archived into a .tar file and then uploaded to COS. If you select "Unarchived files" as the backup file format, the .bak file of each database in the instance will be directly uploaded to COS without being archived.

Backup task execution

The **Execute Backup Task** option allows you to decide whether to back up on the primary or replica node. By default, data is backed up on the primary node.

Viewing and downloading a backup

For easy backup query, management, and analysis, you can [view](#) and [download](#) backups in the console. You can also query the information of all backups or backups for today, the past 7, 15, or 30 days, or a custom period, including backup start/end time, name, policy, mode, file format, size, database, and status.

Viewing the backup space

The backup space occupied by TencentDB for SQL Server instance backup files is allocated by region. It is equivalent to the total storage capacity used by all database backups in a region, including automatic data backups, manual data backups, and log backups.

Increasing the backup retention period or frequency will use more database backup space. You can view the backup

space statistics and trends of all instances in each region under your account as well as the real-time backup space statistics of each instance. For more information, see [Viewing Backup Space](#).

Backup Fees

Last updated : 2024-01-18 17:23:30

TencentDB for SQL Server backup includes local backup and cross-region backup. The former stores the backup files of all TencentDB for SQL Server instances in a local region, while the latter stores backup files in a non-local region.

When you purchase TencentDB for SQL Server primary instances, you will receive free backup capacity. The backup space range is calculated separately for each region of an account. Excess backup files will be charged if your backup space exceeds the free tier. Note that the free tier is not applicable to cross-region backups, and as long as cross-region backup is enabled for your instances, all the generated cross-region backup files will incur fees.

This document describes backup fees. For more information, see:

[Backup Space Billing](#)

[Cross-Region Backup Billing](#)

Local Backup Billing Mode

Backup space is billed in pay-as-you-go mode. Fees will be charged for actual backup space usage beyond the free tier.

Local Backup Pricing

Chinese Mainland	Hong Kong (China) and Other Countries/Regions
0.000113 USD/GB/hour	0.000127 USD/GB/hour

Note:

Billable space of less than 1 GB is not billed, and a billable time period of less than one hour is counted as one hour.

Cross-Region Backup Billing Mode

Cross-region backup fees consist of **storage** and **traffic** fees:

Cross-region backup fees = cross-region backup storage fees + cross-region replication traffic fees

Cross-Region Backup Fees

Storage pricing

Cross-region backup storage fees are charged in a pay-as-you-go (postpaid) manner at the storage price of the destination region as indicated below:

Chinese Mainland and Finance Zones	Hong Kong (China)
0.000113 USD/GB/hour	0.000127 USD/GB/hour

Note:

Billable space of less than 1 GB is not billed, and a billable time period of less than one hour is counted as one hour.

Traffic pricing

Cross-region backup traffic fees are charged based on the linkage between the source and destination regions in a pay-as-you-go (postpaid) manner at the prices as indicated below:

Source Region of Cross-Region Backup	Destination Region of Cross-Region Backup	Traffic Price (USD/GB/Hour)
Chinese mainland (Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, and Nanjing)	Chinese mainland (Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, and Nanjing)	0.09230769
Finance zones (Beijing Finance, Shanghai Finance, and Shenzhen Finance)	Finance zones (Beijing Finance, Shanghai Finance, and Shenzhen Finance)	0.15384615
Hong Kong (China)	Chinese mainland (Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, and Nanjing)	0.35384615
Finance zones (Beijing Finance, Shanghai Finance, and Shenzhen Finance)	Chinese mainland (Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, and Nanjing)	0.12307692
Chinese mainland (Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, and Nanjing)	Finance zones (Beijing Finance, Shanghai Finance, and Shenzhen Finance)	0.12307692

Setting Automatic Backup

Setting Non-Archive Backup Retention

Last updated : 2024-01-18 17:23:30

TencentDB for SQL Server supports automatic backup (including non-archive backup and archive backup) and manual backup.

This document describes how to configure non-archive backup retention in the console.

Note:

The TencentDB for SQL Server instance only backs up data from the business databases rather than the system databases.

The backup is stored in the Beijing time zone (UTC+8) by default. If the default time zone is modified, it will be stored in the time zone of the instance server, but its storage time will still be in Beijing time in the console.

Notes for automatic backup

Backup files are stored in a separate backup space and do not occupy the local disk space of the instance.

Do not perform DDL operations during the backup process to avoid backup failure due to table locking.

Back up your data during off-peak hours.

Backup may take a long time if the data volume is large.

Description of automatic backup

Increasing the retention period of data and log backups may incur additional backup space fees.


Shortening the retention period of log backups may affect the data rollback cycle of the instance.

The free tier for backups is limited. Any excess will be billed on a pay-as-you-go basis.

Local backups are included in the free tier, while cross-region backups are not, that is, cross-region backups incur fees.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Project ▾	Version ▾	Configuration	Network	AZ ▾	Private Netwo Addre:
<input type="checkbox"/> mssql-5jj7 fe0642ad-04dd-4558- b808-61	 Running	Default Project	SQL Server 2008 R2 Enterprise	1-core-2GB/10GB Dual-Server High Availability		Guangzhou Zone 3	172.1 3

3. On the instance management page, select **Backup Management** and click **Auto-Backup Settings**.

4. In the pop-up window, set the following configuration items, read and select **Backup Space Billing Notes**, and click **Save**.

Auto-Backup Settings

Data Backup Settings

Start Time

Data Backup Retention Period day(s)
7-1,830 days. Expired backups will be automatically deleted.

Backup Cycle ☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat ☒ Sun

Next Time 2023-01-12 03:51:15 ⓘ The backup start time in this region is in the UTC+8 time zone.

Periodic Archive ☐ Close
After periodic archive is disabled, no new archive backups will be generated but existing ones won't be deleted immediately. To clear them, please shorten the archive backup retention period.

Log Backup Settings

Log Backup Retention Period day(s)
3-1,830 days. The retention time of log backups must be the same as that of data backups. Expired backup sets are automatically deleted.

Log Backup Frequency Every 10 minutes

Backup Description

1. Prolonging the retention days of data backup and log backup may incur charges for additional backup space.
2. Shortening the retention days of log backup may affect the data rollback period of the instance.
3. Free backup space is limited. Backup space in excess of the free tier will be billed as incurred.

☒ I have read the [Backup Space Billing Notes](#)

Parameter	Description
Start Time	You can set a start time as needed. We recommend that you set it to a value during off-peak hours. This is just the start time of the backup process and does not indicate the end time.

	For example, if the time range is set to 01:00–02:00 AM, the system will initiate a backup at a time point during 01:00–02:00 AM, which depends on the backend backup policy and backup system conditions.
Data Backup Retention Period	It is 7 days by default and can be any value between 7 and 1,830 days. Backup sets will be automatically deleted upon expiration.
Backup Cycle	<p>You can select any time between Monday and Sunday for backup.</p> <p>Note:</p> <p>To ensure the continuity between the data backup retention period and backup cycle:</p> <p>If the backup retention period is set to a value smaller than seven days, the backup cycle will be daily, with Monday through Sunday being selected by default.</p> <p>If the backup retention period is set to a value equal to or greater than seven days, at least two backups will be required every week. If you select only one, you cannot click Save.</p>
Next Time	Set the time when the configured backup cycle will start.
Periodic Archive	The Periodic Archive switch is toggled off when you set non-archive backup retention.
Log Backup Retention Period	It is the same as the data backup retention period by default and cannot be changed. That is, if the data backup retention period is set to 30 days, the log backup retention period will also be 30 days. Backup sets will be automatically deleted upon expiration.
Log Backup Frequency	Backup is performed once every 10 minutes by default.

Setting Archive Backup Retention

Last updated : 2024-01-18 17:23:30

TencentDB for SQL Server supports automatic backup (including non-archive backup and archive backup) and manual backup.

This document describes how to configure archive backup retention in the console.

Notes for automatic backup settings

Backup files are stored in a separate backup space and do not occupy the local disk space of the instance.

Do not perform DDL operations during the backup process to avoid backup failure due to table locking.

Back up your data during off-peak hours.

Backup may take a long time if the data volume is large.

Description of automatic backup settings

Increasing the retention period of data and log backups may cause additional backup space fees.

Shortening the retention period of log backups may affect the data rollback cycle of the instance.

The free tier for backups is limited, and any excess will be billed on a pay-as-you-go basis.

Local backups are included in the free tier, while cross-region backups are not, that is, cross-region backups incur fees.

Periodic archive takes affect only for local backups but not [cross-region backups](#), which will be retained for the period configured in the cross-region backup settings.

Differences between non-archive backup and archive backup

Differences

Archive backup is a more flexible backup policy on the basis of non-archive automatic backup. It supports setting the number of retained backups by month, quarter, or year and does not need to retain additional new backups. However, its retention period is different from (longer than) that of non-archive backup.


Example

If the non-archive backup cycle in the automatic backup settings of instance A is Monday, Tuesday, and Wednesday, the archive backup retention policy is monthly, and the number of backups to be retained is 2, then the instance will still be automatically backed up every Monday, Tuesday, and Wednesday in the month, but the backup retention period will vary by archive backup policy. Among these backups, non-archive backups will be retained for the configured data backup retention period, while archive backups will be retained for the configured archive backup retention period. For example, if the system performs archive backup on a certain Wednesday, then the backup generated on the day will be retained for the archive backup retention period rather than in two copies for different retention periods.

</dx-alert>

Enabling archive backup retention

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.

<input type="checkbox"/>	Instance ID/Name	Status ▾	Project ▾	Version ▾	Configuration	Network	AZ ▾	Private Netwo Addre:
<input type="checkbox"/>	mssql-5jj7 fe0642ad-04dd-4558- b808-61	 Running	Default Project	SQL Server 2008 R2 Enterprise	1-core-2GB/10GB Dual-Server High Availability		Guangzhou Zone 3	172.1 3

3. On the instance management page, select **Backup Management** and click **Auto-Backup Settings**.
4. In the pop-up window, set the following configuration items, read and select **Backup Space Billing Notes**, and click **Save**.

Auto-Backup Settings

Data Backup Settings

Start Time

03:00 - 04:00

Data Backup Retention Period

7

day(s)

7-1,830 days. Expired backups will be automatically deleted.

Backup Cycle

☒ Mon
 ☒ Tue
 ☒ Wed
 ☒ Thu
 ☒ Fri
 ☒ Sat
 ☒ Sun

Next Time

2023-01-12 03:51:15 ⓘ The backup start time in this region is in the UTC+8 time zone.

Periodic Archive

☒ Enable

Archive Backup Retention Period

—

365

+

day(s)

90-3650 days. Expired backup sets will be automatically deleted.

Archive Backup Retention Policy

Monthly

1

pcs ⓘ Retention Policy

1-{{num}}

Start Time

2023-01-11

View Retention Plan

Log Backup Settings

Log Backup Retention Period

7

day(s)

3-1,830 days. The retention time of log backups must be the same as that of data backups. Expired backup sets are automatically deleted.

Log Backup Frequency

Every 10 minutes

Backup Description

1. Prolonging the retention days of data backup and log backup may incur charges for additional backup space.

2. Shortening the retention days of log backup may affect the data rollback period of the instance.

3. Free backup space is limited. Backup space in excess of the free tier will be billed as incurred.

☒ I have read the [Backup Space Billing Notes](#)

Save

Cancel

Parameter	Description
Start Time	You can set a start time as needed. We recommend that you set it to a value during off-peak hours. This is just the start time of the backup process and does not indicate the end time. For example, if the time range is set to 01:00–02:00 AM, the system will initiate a backup at a time point during 01:00–02:00 AM, which depends on the backend backup policy and backup system conditions.
Data Backup	It is 7 days by default and customizable between 7 and 1,830 days. Expired backup sets will be automatically deleted.

Retention Period	
Backup Cycle	<p>You can select any time between Monday and Sunday for backup.</p> <p>Note: To ensure the continuity between the data backup retention period and backup cycle: If the backup retention period is set to a value smaller than 7 days, the backup cycle will be daily, with Monday through Sunday being selected by default. If the backup retention period is set to a value equal to or greater than 7 days, at least two backups will be required every week. If you select only one, you cannot click Save.</p>
Next Time	Set the time when the configured backup cycle will start.
Periodic Archive	Toggle on Periodic Archive and configure backup retention as follows.
Archive Backup Retention Period	<p>It is 365 days by default and customizable between 90 and 3,650 days. Expired backup sets will be automatically deleted.</p> <p>Note: The archive backup retention period should be longer than the non-archive backup retention period.</p>
Archive Backup Retention Policy	<p>You can set the number of retained backups by month, quarter, or year.</p> <p>Note: The maximum number of retained backup varies by time range and should not be exceeded. If the number is 1, the first valid backup in the current cycle will be retained. If the number is 2, the first and middle valid backups in the current cycle will be retained. If the number is above 2, valid backups in the current cycle will be retained by the average time interval.</p>
Start Date	The time to start archive backup.
Log Backup Retention Period	It is the same as the non-archive data backup retention period by default and cannot be changed. Expired backup sets will be automatically deleted.
Log Backup Frequency	Backup is performed once every 10 minutes by default.

Modifying periodic archive

After periodic archive is enabled, you can modify the retention policy by setting parameters as instructed in [Enabling periodic archive](#).

Note:

Modifying the archive backup retention period will affect the retention period of historical archive backups and cause expired backups to be cleared.

Modifying the archive backup retention policy will affect only new but not previously retained archive backups.

Modifying the start time will affect only new but not previously retained archive backups.

Disabling period archive

Note:

After it is disabled, no new archive backups will be generated.

After it is disabled, existing archive backups will be retained and automatically deleted upon expiration according to the original policy. You can also clear them by modifying the archive backup retention period.

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.
3. On the instance management page, select **Backup Management** and click **Auto-Backup Settings**.
4. Toggle off **Periodic Archive**.

Auto-Backup Settings

Data Backup Settings

Start Time

03:00 - 04:00

Data Backup Retention Period

7

day(s)

7-1,830 days. Expired backups will be automatically deleted.

Backup Cycle

☒ Mon ☒ Tue ☒ Wed ☒ Thu ☒ Fri ☒ Sat ☒ Sun

Next Time

2023-01-12 03:51:15 ⓘ The backup start time in this region is in the UTC+8 time zone.

Periodic Archive

☐ Close

After periodic archive is disabled, no new archive backups will be generated but existing ones won't be deleted immediately. To clear them, please shorten the archive backup retention period.

Log Backup Settings

Log Backup Retention Period

7

day(s)

3-1,830 days. The retention time of log backups must be the same as that of data backups. Expired backup sets are automatically deleted.

Log Backup Frequency

Every 10 minutes

Backup Description

1. Prolonging the retention days of data backup and log backup may incur charges for additional backup space.

2. Shortening the retention days of log backup may affect the data rollback period of the instance.

3. Free backup space is limited. Backup space in excess of the free tier will be billed as incurred.

☒ I have read the [Backup Space Billing Notes](#)

Save

Cancel

Viewing the retention plan

After setting the archive backup retention policy and start time, you can click **View Retention Plan** next to the start time to preview the backup plan.

起始日期

2022-01-11

📅

查看保留计划

Blue dates are for non-archive backups.

Red dates are for archive backups.

You can click **Non-archive Backup** or **Archive Backup** to hide corresponding dates for easier preview.

The backup plan preview is currently for backups for the year to come and is for reference only.

Example 1: Retaining one backup per month starting from December 1, 2022, with the backup cycle being Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, and Sunday.

Example 2: Retaining one backup per month starting from November 22, 2022, with the backup cycle being Monday, Wednesday, Friday, and Sunday.

Example 3: Retaining four backups per quarter starting from November 22, 2022, with the backup cycle being Monday, Wednesday, and Friday.

Example 4: Showing only dates in the archive backup retention plan.

Creating Manual Backup

Last updated : 2024-07-30 16:34:32

TencentDB for SQL Server supports automatic backup and manual backup. This document describes how to create a manual backup in the console.

Note:

Backup files are stored in a separate backup space and do not occupy the local disk space of the instance.

Do not perform DDL operations during the backup process to avoid backup failure caused by table locking.

Back up your data during off-peak hours.

Backup may take a long time if the data volume is large.

We recommend you not modify the restoration mode for Basic Edition instances, which is **full** mode by default. If you modify it to the **simple** mode, the log chain will be broken and manual backup will fail.

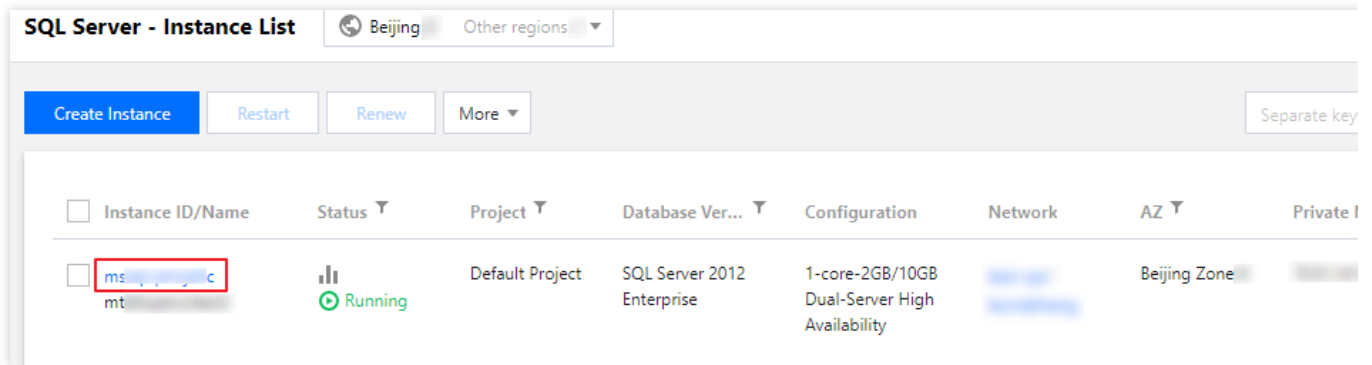
The backup link is as follows: Initially, backup files are uploaded to a local temporary backup space before being transferred to the COS backup space. Backup files within the local temporary backup space are retained for one day. For the cloud disk edition architecture, the size of the local temporary backup space is 50% of the purchased disk size. In instances where the backup files are considerably large, there may arise a scenario where the backup process is hindered due to insufficient local temporary backup space.


Prerequisites

Before performing manual backup, make sure that you have created a database. For more information, see [Creating Database](#).

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



Beijing		Other regions					
Create Instance	Restart	Renew	More	Separate key			
<input type="checkbox"/> Instance ID/Name	Status	Project	Database Ver...	Configuration	Network	AZ	Private IP
<input type="checkbox"/> ms-sql-12345678 mt	 Running	Default Project	SQL Server 2012 Enterprise	1-core-2GB/10GB Dual-Server High Availability	VPC	Beijing Zone	192.168.1.1

3. On the **Backup Management** tab, click **Manual Backup**.

4. In the pop-up window, enter the backup name, select the backup policy, and click **OK**.

Backup Name: Set the manual backup name, which can contain up to 128 letters, digits, or underscores.

Backup Policy:

Instance backup: Backs up all databases in the entire instance.

Multi-database backup: Backs up selected databases.

Note:

The time it takes to complete the backup is subject to the backup file size and ranges from around 5 to 120 minutes.

5. After the manual backup is completed, you can query it in the backup list on the **Backup Management** tab.

Setting Backup Task

Last updated : 2024-01-18 17:23:30

Backup task settings include global variables of manual and scheduled backup, such as the backup file format and backup task execution option.


This document describes how to configure a backup task in the console.

Note:

Only two-node 2017/2019 instances support configuring backup task execution options.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the **Backup Management** tab, click **Backup Task Settings**.
4. Specify configuration items in the pop-up window and click **Save**.

Upload Backup File: You can specify the format of backup files (**Archive file** or **Unarchived files**). By default, backup files (i.e., .bak files) will be archived into a .tar file and then uploaded to COS. If you select **Unarchived files** as the backup file format, the .bak file of each database in the instance will be directly uploaded to COS without being archived.

Execute Backup Task: You can decide whether to back up on the primary or replica node. By default, data is backed up on the primary node.

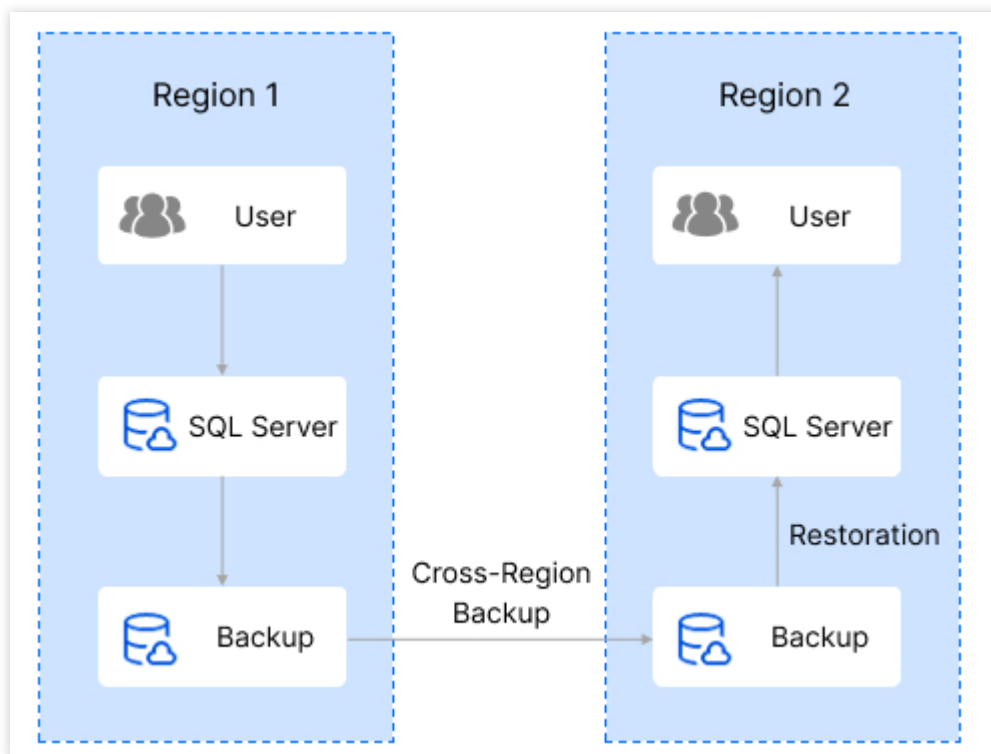
Cross-Region Backup

Last updated : 2024-01-18 17:23:30

This document describes the cross-region backup feature of TencentDB for SQL Server. This feature allows you to store backup files in another region, which enhances the regulatory compliance and disaster recovery capabilities and improves the data reliability.

Overview

Data is an important part of enterprise operations. Although the information technology brings convenience, it also reveals that electronic data and stored information are very vulnerable to damage or loss. Any incident, such as natural disaster, system failure, maloperation, or virus, can cause interruption of business operations or even disastrous losses. Therefore, ensuring the security and integrity of core data is a top priority of every enterprise. The cross-region backup feature of TencentDB for SQL Server can be used to store backup files in another region so as to minimize data corruptions caused by natural disasters or system failures. This feature ensures the high availability, security, and recoverability of data and implement various features, such as remote backup and restoration, remote disaster recovery, long-term data archive, and regulatory compliance.



Notes on cross-region backup

Cross-region backup doesn't affect the local default backup, and both coexist after cross-region backup is enabled. Cross-region backup will be triggered after the local default automatic backup is complete, that is, the default automatic backup is dumped to the storage device for cross-region backup.

Backup files in the cross-region backup space include automatic data backups and log backups, that is, local automatic backups are automatically synced to the destination region for storage.

The retention period of cross-region backups is 7 days by default and customizable between 7 and 1,830 days.

Enabling/Disabling cross-region backup or changing the backup region won't affect existing backups.

The retention period of cross-region backups only affects their lifecycle.

Modifying the cross-region backup retention period will affect the lifecycle of existing cross-region backups.

Changing the backup region option for cross-region backup won't affect the backup and storage regions of existing cross-region backups.

Cross-region backups stored in another region cannot use the free storage space provided for local backups.

Billing

Cross-region backup fees consist of **storage** and **traffic** fees:

Cross-region backup fees = cross-region backup storage fees + cross-region replication traffic fees

Cross-region backup storage and traffic fees are charged based on the destination region and the linkage between the source and destination regions respectively in a pay-as-you-go manner.

For billing details, see [Cross-Region Backup Billing](#).

Differences between cross-region backup and local backup

Item	Local Backup	Cross-Region Backup
Enabled by default	Yes.	No. It needs to be enabled manually.
Backup storage region	Instance region.	Destination region.
Region for backup billing	The backup volume is counted in the backup space for the backup region of the primary instance.	The backup volume is counted in the backup space for the backup region of the primary instance, and backup fees are calculated at the

		price of the backup storage region (i.e., the destination region).
Billing cycle	For a pay-as-you-go instance, local backups are retained for three days by default after the instance is moved to the recycle bin and are terminated along with the instance after three days.	For a pay-as-you-go instance, local backups are retained for three days by default after the instance is moved to the recycle bin and are terminated along with the instance after three days.
Uses the free storage space	Yes.	No.
Backup space billing rules	Before the backup size exceeds the free tier, backups are free of charge. After the backup size exceeds the free tier, the hourly backup fees = (total size of backup files - free tier) * backup unit price. Billable space = data backup volume (automatic + manual) + log backup volume (automatic) - free backup space (all values are for the current region).	After cross-region backup is enabled, all the generated cross-region backup files will incur fees. Billable space (for the region of the primary instance) = data backup volume (automatic) + log backup volume (automatic) (all values are for the destination region).

Supported regions

Source Region	Supported Destination Region
Beijing, Guangzhou, Shanghai, Chengdu, Chongqing, Nanjing, Beijing Finance, Shanghai Finance, Shenzhen Finance, and Hong Kong (China)	Beijing, Guangzhou, Shanghai, Chengdu, Chongqing, Nanjing, Beijing Finance, Shanghai Finance, and Shenzhen Finance. Any region other than the source region, subject to the supported regions as indicated in the console.


Note:

Currently, you cannot perform a cross-region backup task from a Chinese mainland region to Hong Kong (China).
Currently, cross-region backup is not supported outside the Chinese mainland but will be supported in more regions in the future.

Enabling cross-region backup in the console

1. Log in to the [TencentDB for SQL Server console](#).

- In the instance list, click the ID or **Manage** in the **Operation** column of the target instance to enter the **Instance Management** page.
- On the **Instance Management** page, select **Backup Management** and click **Cross-Region Backup Settings**.



- In the **Cross-Region Backup Settings** window, set the following parameters and click **Save**.

Cross-Region Backup Settings

i Cross-region backup is used to ensure the high availability and security of your data. when it is enabled, the local backup will not be affected, and both backups will exist at the same time. For details, see [Cross-Region Backup Feature Notes](#). Free local storage space is unavailable for cross-region backup. For details, see [Cross-Region Backup Billing Notes](#).

Cross-Region Backup ☒

Backup Region Beijing
Select additional storage region

Cross-Region Backup Retention Period 7 day(s)
7-1,830 days. Expired backups will be automatically deleted.

☐ Enabling cross-region backup will incur additional charges. I have read [Cross-Region Backup Billing Notes](#).

Save Cancel

Parameter	Description
Cross-Region Backup	Switch for the cross-region backup feature, which can be toggled on.
Backup Region	The region where to store cross-region backups. Click the drop-down list to select one or two target regions.
Cross-Region Backup Retention Period	The period for retaining cross-region backup files, which can be set between 7 and 1,830 days.
Cross-Region Backup	Enabling cross-region backup will incur additional charges. Read the notes and

Billing Notes

indicate your consent.

Modifying cross-region backup settings

Note:

Changing the backup region will affect the storage region of new backups only but not existing backups.

Modifying the cross-region backup retention period will affect the lifecycle of both existing and new cross-region backups.

Option 1

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click the ID or **Manage** in the **Operation** column of the target instance to enter the **Instance Management** page.
3. On the **Instance Management** page, select **Backup Management** and click **Cross-Region Backup Settings**.
4. In the **Cross-Region Backup Settings** window, modify the backup region and cross-region backup retention period and click **Save**.

Option 2

1. Log in to the [TencentDB for SQL Server console](#).
2. Select **Database Backup** on the left sidebar, and select a region at the top of the **Database Backup** page.
3. On the **Database Backup** page, select **Overview** > **Real-Time Backup Statistics** > **Cross-Region Backup**.
4. In the cross-region backup list, click **Cross-Region Backup Settings** in the **Operation** column of the target instance.
5. In the **Cross-Region Backup Settings** window, modify the backup region and cross-region backup retention period and click **Save**.

Disabling cross-region backup

Note:

After cross-region backup is disabled, no cross-region backup files will be generated and billed.

Option 1

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, click the ID or **Manage** in the **Operation** column of the target instance to enter the **Instance Management** page.
3. On the **Instance Management** page, select **Backup Management** and click **Cross-Region Backup Settings**.
4. In the **Cross-Region Backup Settings** window, toggle off the **Cross-Region Backup** feature.
5. In the pop-up window, select **Retain** or **Delete** and click **OK**.

Note:

If **Retain** is selected, existing cross-region backup files will be retained and billed for the configured cross-region backup retention period, which can be modified before cross-region backup is disabled.

If **Delete** is selected, existing cross-region backup files will be deleted when cross-region backup is disabled.

Option 2

1. Log in to the [TencentDB for SQL Server console](#).
2. Select **Database Backup** on the left sidebar, and select a region at the top of the **Database Backup** page.
3. On the **Database Backup** page, select **Overview > Real-Time Backup Statistics > Cross-Region Backup**.
4. In the cross-region backup list, click **Cross-Region Backup Settings** in the **Operation** column of the target instance.
5. In the **Cross-Region Backup Settings** window, toggle off the **Cross-Region Backup** feature.
6. In the pop-up window, select **Retain** or **Delete** and click **OK**.

Viewing Backup List

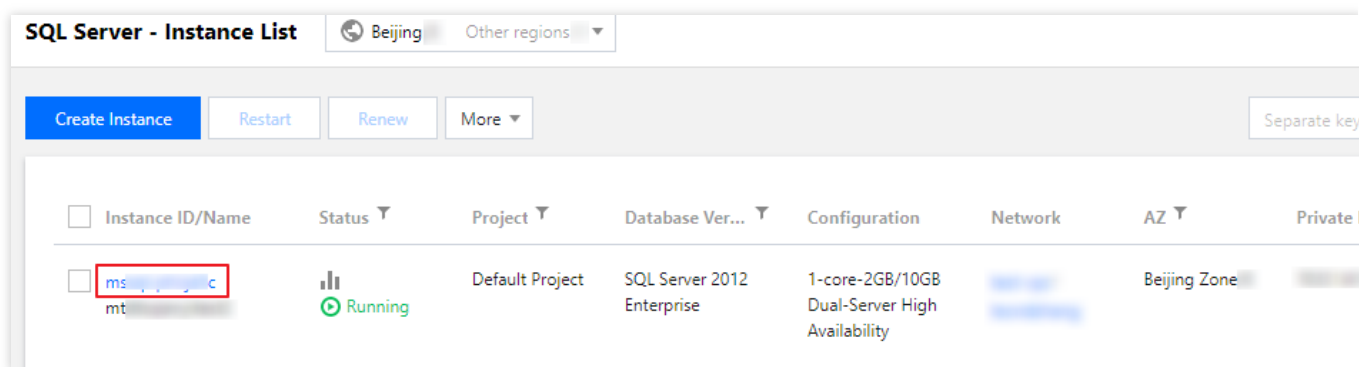
Last updated : 2024-01-18 17:23:30

TencentDB for SQL Server automatically backs up data based on the default backup settings. You can modify the automatic or manual backup settings in the console. You can also view the backup files and relevant information in the backup list.

This document describes how to view backup files in the console.

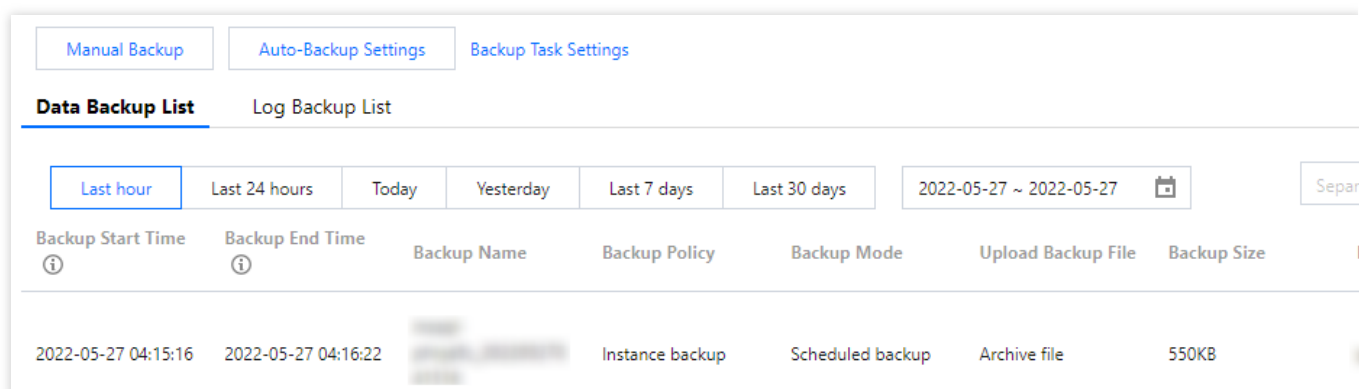
Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select the **Backup Management** tab, view backup tasks in the data or log backup list.

The data backup list displays the backup start/end times, name, policy, mode, file format, and size, database, region, status, and operations (download, restore, delete).



The log backup list displays the file name, log data start/end times, backup size, region, and operation (download).

Manual Backup

Auto-Backup Settings

Backup Task Settings

Data Backup List

Log Backup List

Last hour

Last 24 hours


Today

Yesterday

Last 7 days

Last 30 days

2022-05-27 ~ 2022-05-27



File Name	Log Backup Start Time ⓘ	Log Backup End Time ⓘ	Backup Size
2022-05-27 14:54:15	2022-05-27 14:54:15	2022-05-27 14:57:06	80K
2022-05-27 14:24:16	2022-05-27 14:24:16	2022-05-27 14:27:06	80K

You can filter backup files by time (last hour, last 24 hours, today, yesterday, last 7 days, last 30 days, or custom time period).

Manual Backup

Auto-Backup Settings

Backup Task Settings

Data Backup List

Log Backup List

Last hour

Last 24 hours

Today

Yesterday

Last 7 days

Last 30 days

2022-05-

File Name

Log Backup Start Time ⓘ

Log Backup End Time ⓘ

You can search for backup files in the search box on the right by database name, backup name, policy, mode, or file format.

Separate keywords with " "; press Enter to separate filter tags		
Select a filter	Status	Operation
Backup Name		
Database Name		
Backup Policy	Successful	Download Restore
Backup Mode		Delete
Upload Backup File		

FAQs

1. Can I download or restore backup files that exceed the retention period?

Expired backup sets will be deleted automatically and cannot be downloaded or restored.

We recommend that you configure a backup retention period based on business needs or download the backup files locally in the [TencentDB for SQL Server console](#).

You can also manually back up the instance data in the console.

Note:

Manual backups will also take up the backup space. We recommend that you plan the usage of the backup space appropriately to reduce costs.

2. Can I delete backups manually?

Automatic backups cannot be deleted manually. You can set the retention period for automatic backups, and they will be deleted automatically upon expiration. If [Cross-Region Backup](#) is enabled, you can also set the retention period for cross-region backup files, which will be deleted automatically upon expiration.

Manual backups can be manually deleted from the backup list in the [TencentDB for SQL Server console](#). If they are not manually deleted, they will be retained for the same period as automatic backups.

3. Can I disable data and log backups?

No. However, you can lower the backup frequency and delete manual backup files that are no longer used in the [TencentDB for SQL Server console](#) to reduce the backup space usage.

4. How can I reduce the backup capacity costs?

Delete manual backups that are no longer used. You can log in to the [TencentDB for SQL Server console](#), click an instance ID/name to access the instance management page, and delete manual backups on the **Backup Management** tab.

Reduce the frequency of automatic data backup for non-core businesses. You can adjust the backup cycle and backup file retention period in the console, which should be at least twice a week.

Note:

The rollback feature relies on the backup cycle and retention days of data backups and log backups. Rollback will be affected if you reduce the automatic backup frequency and retention period. You can select the parameters as needed. For more information, see [Rolling back Databases](#).

Reduce the retention period of data and log backups for non-core businesses (a 7-day retention period can meet the requirements of most scenarios).

Business Scenario	Recommended Backup Retention Period
Core businesses	7-1,830 days
Non-core, non-data businesses	7 days
Archival businesses	7 days. We recommend you manually back up data based on your actual business

	needs and delete the backups promptly after use
Testing businesses	7 days. We recommend you manually back up data based on your actual business needs and delete the backups promptly after use

5. Can I download the backup files of an isolated instance?

Yes.

A pay-as-you-go instance will be isolated and moved into the recycle bin 24 hours after expiration. At this time, rollback and manual backup will be prohibited, but automatic backup can still be downloaded by clicking **More** in the **Operation** column of the instance. Excessive backup space of the instance will continue to be billed until the instance is eliminated.

Downloading Backup

Last updated : 2024-07-31 12:41:20

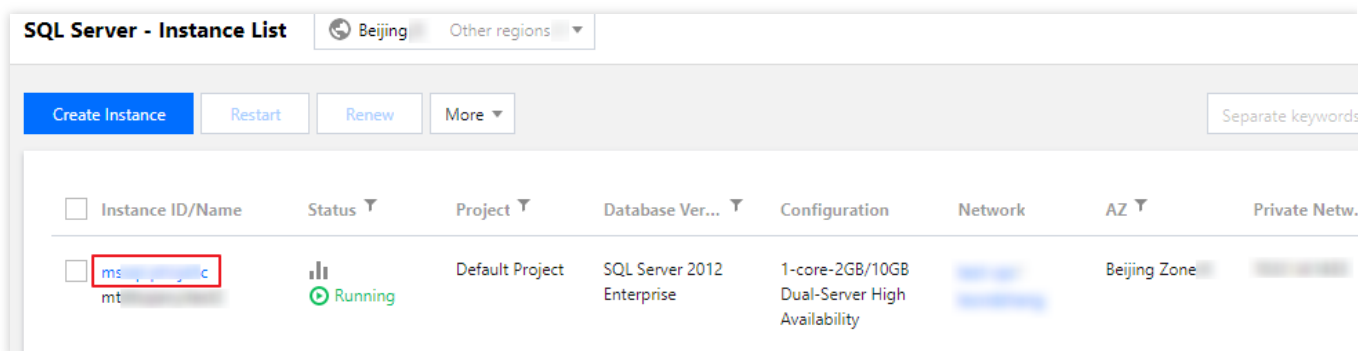
The TencentDB for SQL Server console provides the list of backup files that can be downloaded over the private or public network. You can use the downloaded backups to restore data from one database to another (such as a self-built one).

This document describes how to download a local or cross-region backup in the console.

Directions

Downloading the backup file of a running instance

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select **Backup Management** and view the target backup file in the data or log backup list.

To download an archive backup file, click **Download** in the **Operation** column in the backup list. The downloaded backup is a local backup if its region is the same as that of the instance. If not, the downloaded backup is a cross-region backup.

Instance Details

System Monitoring

Account Management

Database Management

Backup Management

Slow log

Backup List

Backup and log settings

Creating Backups

Start/End time of backup	Backup Policy	Backup Method	Backup Size
2019-07-31 20:01:04/2019-07-31 20:02:57	Multi-database backup	Scheduled backup	452KB
2019-07-31 19:56:52/2019-07-31 19:58:55	Multi-database backup	Scheduled backup	452KB

To download an unarchived backup file, select **View Details > Download** in the **Operation** column in the backup list to download the backup file of each database.

Manual Backup

Auto-Backup Settings

Backup Task Settings

Data Backup List

Log Backup List

Last hour

Last 24 hours

Today

Yesterday

Last 7 days

Last 30 days

2022-05-27 ~ 2022-05-27

Separate

Backup Start Time	Backup End Time	Backup Name	Backup Policy	Backup Mode	Upload Backup File	Backup Size	Data
2022-05-27 00:27:15	2022-05-27 00:28:08	Instance backup	Scheduled backup	Unarchived files	384KB		

4. In the pop-up window, get the file download address for fast download over the private network by running the `wget` command, or directly click **Download**.

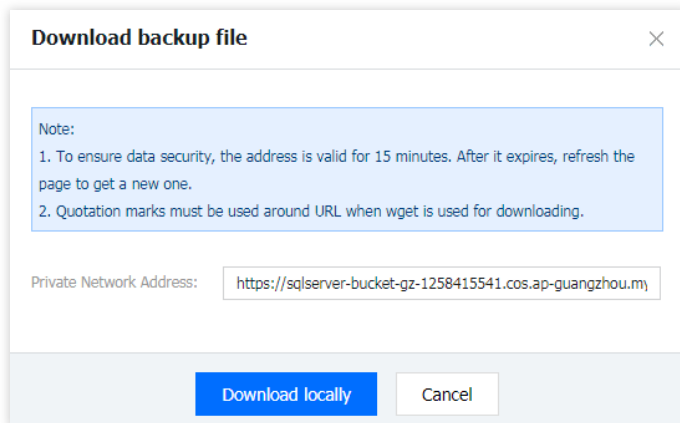
Note:

We recommend that you copy the private download link, log in to a Linux CVM instance in the same VPC as the TencentDB instance, and run the `wget` command for download over the private network at a higher speed. For more information, see [Customizing Linux CVM Configurations](#).

The download address is valid for 15 minutes, after which you will need to enter the download page again to get a new one.

When a backup file is downloaded by using `wget` via the private network, the command format is `wget -c 'backup file download URL' -O 'custom filename'`, with the addition of single quotes in English for the URL required.

When a backup file is downloaded by using `wget` via the public network, the command format is `wget -c "backup file download URL" -O 'custom filename'`, with the addition of double quotes in English for the URL required.



Downloading the backup file of an isolated instance

You can also download the backup file of an isolated instance.

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click **More > Download Backup** in the **Operation** column to enter the backup download page.

Create Instance Restart Renew More ▾ Separate keyword							
<input type="checkbox"/> Instance ID/Name	Status ▾	Project ▾	Database Ver... ▾	Configuration	Network	AZ ▾	Private Netw
<input type="checkbox"/> n- n-	Running	Default Project	SQL Server 2012 Enterprise	1-core-2GB/10GB Dual-Server High Availability		Beijing Zone 6	
<input type="checkbox"/> ms- mt-	Isolated	Default Project	SQL Server 2012 Enterprise	1-core-2GB/20GB Dual-Server High Availability		Beijing Zone 5	
<input type="checkbox"/> m- 2- b-	Running	Default Project	SQL Server 2017 Enterprise	1-core-2GB/10GB Cluster Edition		Beijing Zone 5	

To download an archived backup file, click **Download** in the **Operation** column in the backup list, and download the file on the displayed download page.

To download unarchived backup files, select **View Details > Download** in the **Operation** column to download the backup file of each database.

Manual Backup

Auto-Backup Settings

Backup Task Settings

Data Backup List

Log Backup List

Last hour

Last 24 hours

Today

Yesterday

Last 7 days

Last 30 days

2022-05-27 ~ 2022-05-27

Separate key

Backup Start Time	Backup End Time	Backup Name	Backup Policy	Backup Mode	Upload Backup File	Backup Size	Databa
2022-05-27 00:27:15	2022-05-27 00:28:08		Instance backup	Scheduled backup	Unarchived files	384KB	

Deleting Manual Backup

Last updated : 2024-01-18 17:23:30

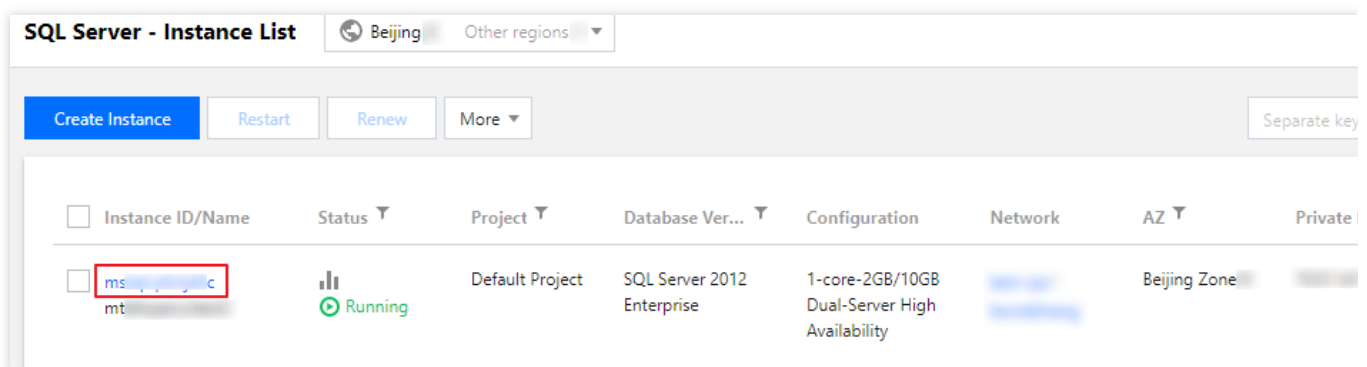
This document describes how to delete manual backup files in the console to reduce the backup space costs. Automatic backup files cannot be manually deleted but will be automatically deleted upon expiration.

Note:

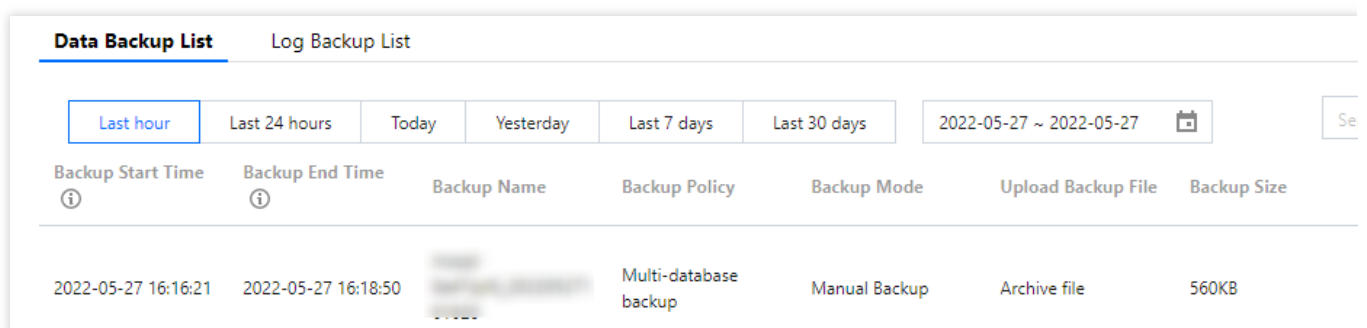
Manual backups cannot be restored once deleted.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the target instance ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select the **Backup Management** tab and click **Delete** in the **Operation** column in the backup list.



4. In the pop-up window, confirm the backup file to be deleted and click **OK**.

Note:

You can batch delete historical manual backup files. To do so, click **Batch Delete Manual Backups**, select the time range of the files to be deleted in the settings window, and click **OK**.

Viewing Backup Space

Last updated : 2024-01-18 17:23:30

This document describes how to view the backup space in the console.

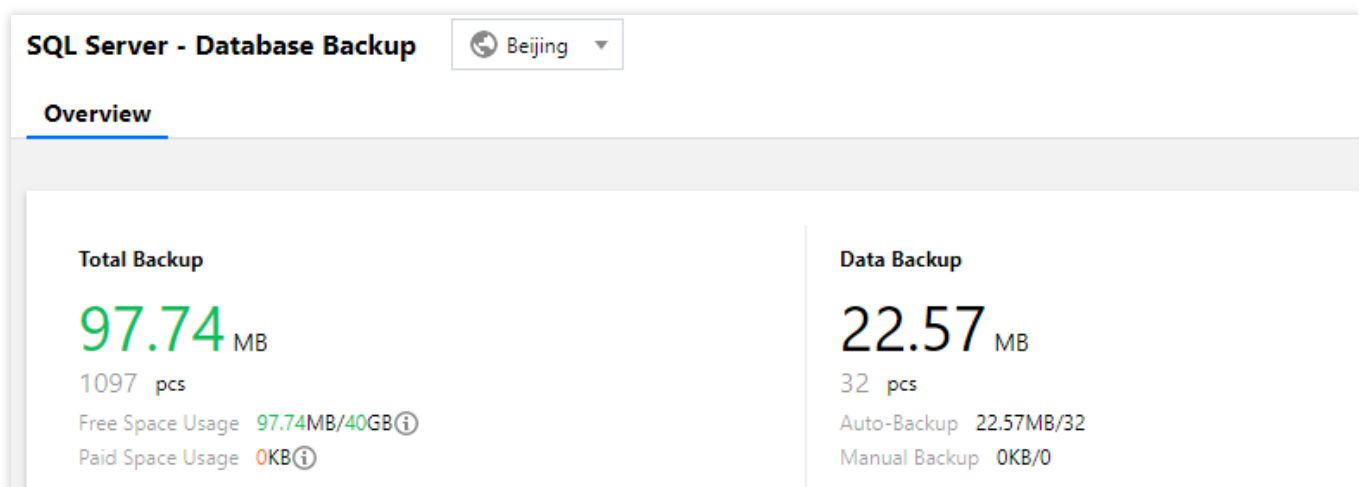
Overview

The backup space occupied by TencentDB for SQL Server instance backup files is allocated by region. It is equivalent to the total storage capacity used by all database backups in a region, including local and cross-region automatic data backups, manual data backups, and log backups. For cross-region backup, the backup volume is counted in the backup space for the backup region of the primary instance, and backup fees are calculated at the price of the backup storage region (i.e., the destination region).

Increasing the backup retention period or frequency will use more database backup space. You can view the backup space statistics and trends of all instances in each region under your account as well as the real-time backup space statistics of each instance.

Viewing the total backup

1. Log in to the [TencentDB for SQL Server console](#).
2. On the left sidebar, select **Database Backup**.
3. On the **Database Backup** page, select a region to view the backup space and number of all backups, data backups, and log backups on the **Overview** tab, including those of local backup and cross-region backup.




Note:

In the **Total Backup** section, you can view the usage details of the free space. Cross-region backup cannot use the free space; in other words, the space used by the backup files generated by cross-region backup will incur fees.

Viewing the backup trend

1. Log in to the [TencentDB for SQL Server console](#).
2. On the left sidebar, select **Database Backup**.
3. On the **Database Backup** page, select a region. Then, you can view the trend of total local or cross-region backup space size at the bottom of the **Overview** tab by time (last hour, last 24 hours, today, last 7 days, last 30 days, or a selected custom time period).

Viewing the trend of local backup



Viewing the trend of cross-region backup



Viewing real-time backup statistics

1. Log in to the [TencentDB for SQL Server console](#).
2. On the left sidebar, select **Database Backup**.
3. On the **Database Backup** page, select a region to view the real-time backup statistics at the bottom of the **Overview** tab. For local backup, you can view the instance ID/name, backup space, data backup, log backup, automatic backup, and manual backup. For cross-region backup, you can view the instance ID/name, cross-region backup status, last backup region, backup retention period (day), last backup start time, total backup, data backup, log backup, and operations (cross-region backup settings).

Viewing the real-time statistics of local backup

Real-Time Backup Statistics ⓘ

Local Backup Cross-Region Backup

Separate keywords

<input type="checkbox"/> Instance ID/Name	Backup Space ⚙	Data Backup ⚙	Log Backup ⚙	Auto-I
<input type="checkbox"/> [blurred]	[blurred]	1.46MB 7	29.21MB 997	30.67M 1004

Viewing the real-time statistics of cross-region backup

Real-Time Backup Statistics ⓘ

Local Backup **Cross-Region Backup**

Separate keywords with

<input type="checkbox"/> Instance ID/Name	Cross-Region Back...	Last Backup Region	Backup Retention P...	Last Backup Start T...	Total Backup ⬆	Data Ba
<input type="checkbox"/> <div>...</div>	Enable	...	7	...	OKB 0	OKB 0

Note:

The list of real-time cross-region backup statistics displays only the information of instances with cross-region backup enabled as well as instances with cross-region backup disabled but with existing cross-region backup files retained.

4. On the backup statistics page, you can search for an instance by instance ID to view its real-time backup statistics. You can also perform various operations such as sorting statistical fields by space size, refreshing the statistics page, and downloading data.

Local Backup Cross-Region Backup

Separate keywords with

<input type="checkbox"/> Instance ID/Name	Backup Space ⬆	Data Backup ⬆	Log Backup ⬆	Auto-Back
---	----------------	---------------	--------------	-----------

Rolling back

Rollback Scheme Overview

Last updated : 2024-01-18 17:23:30

This document provides an overview of the data recovery scheme for SQL Server.

Scenario	Feature		Related Operations
Rollback (Recovery of cloud-based databases)	Rollback to the current instance (supporting rollback to the current instance either by time point or backup set)		Rollback to current instance
	Rollback to the existing instance	In-region rollback (Supports rollback to other instances within the same region based on time point or backup set)	Rollback to existing instance - in-region rollback
		Cross-region rollback (Supporting rollback to other instances across regions based on time point or backup set)	Rollback to existing instance (cross-region)
Restoration (Migrating an on-premises database to the cloud)	Migration/restoration to an existing instance		Cold backup migration (Migrating data to TencentDB for SQL Server from another cloud provider or a self-built SQL Server via backup files) Migrating data with DTS (Using DTS to migrate data from another cloud provider or self-built SQL Server to TencentDB for SQL Server) Using DTS for cross-account data migration (Employing DTS to perform data migrations between Tencent SQL Server instances under different accounts)

Rolling back to the current instance

Last updated : 2024-01-18 17:23:30

SQL Server offers a rollback tool to perform rollback operations on instances, and reconstructs historical data through regular backups and real-time transactions. This document describes how to roll back to the current instances for single database, multiple databases, or all the databases according to a time point or a backup set.

Prerequisites

To utilize the rollback feature, there must be SQL Server instances. For instructions on the creation of instances, please refer to [Creating SQL Server Instance](#).

In order to use the rollback feature, a database must exist under an instance. For instructions on the creation of a database, see [Creating a Database](#).


Points of Attention

Before initiating a rollback, please ensure there is sufficient storage space for an instance.

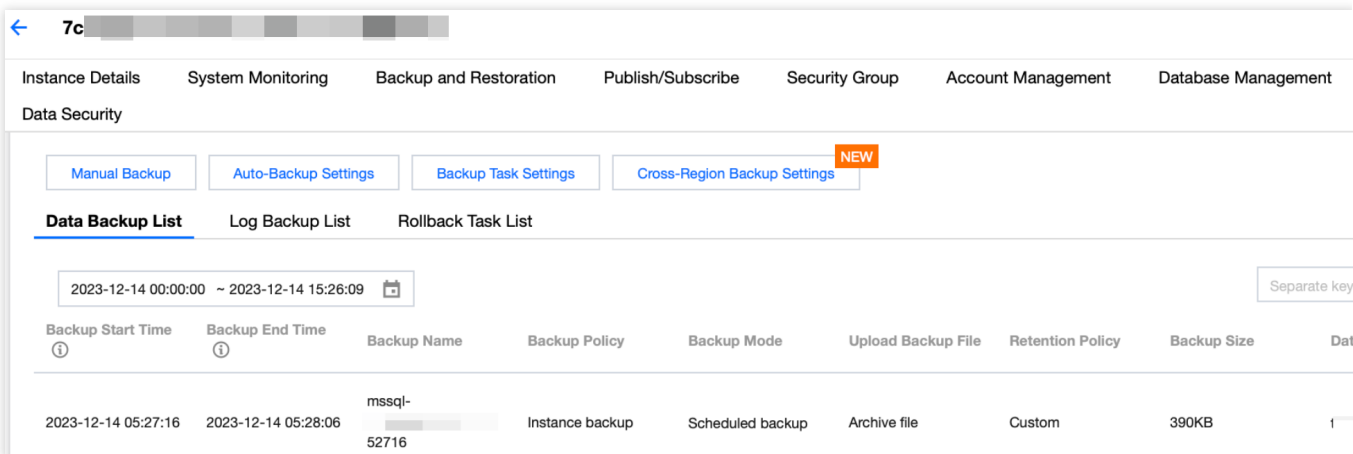
Ensure that no other tasks are being performed for the source instance before initiating a rollback.

The steps are as follows:

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the **region** on the top, locate the required instance, and click **Instance ID** or **Manage** in the **Operation** column to go to the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ
<input type="checkbox"/> mssql- ec-79d- b1-83f	 Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated	Default-VPC - Default-Subnet	Guangzhou Zone 6 Primary Replica	1433

3. Select the **Backup Management** tab, and then either click **Rollback** in the upper right corner of the backup management page or click **Rollback** in the **Data Backup List** operation column.



4. In the pop-up rollback settings interface, complete the corresponding configurations based on the selected rollback method and click **Save**.

Scenario One: Rollback by Time Point

Roll back database to the specified instance

Select the target instance

Rollback to Current instance

Rollback to Existing instance

Select rollback mode

Rollback by Time Point

Rollback by Backup Set

You can roll back the database to a specific time point from a source database instance. The time range depends on the log backup retention period you set.

Set rollback time

2023-12-14 15:27:35

The rollback will be performed in Beijing time zone, with the time range available from 2023-12-07 15:27:45 to 2023-12-14 15:27:45. Please select a time point in this range.

Overwrite original database

☐

Yes

Select whether the database to be rolled back will overwrite the original database.
If you select "Yes", the original database will be renamed to "RESTORE_OLD_*_original database name" when the rollback succeeded, with the old name used for the new database.
If you don't select "Yes", the name of the original database will remain unchanged, and the new database will be named automatically by system and can be renamed later.

Select the database to be rolled back

Select rollback database

Select rollback database

Database Name

Da... Set Database Name After R...

Hold the Shift key down to select multiple items

Save

Cancel

Parameter	Description
Select the target instance	Select Rollback to Current instance .
Select rollback mode	Select Rollback by Time Point . You can roll back the database to a specific time point from a source database instance. The time range depends on the log backup retention period you set.
Set rollback time	Select the rollback time. A uniform rollback time can only be set for databases for the same instance.

Overwrite original database	<p>Select whether the database to be rolled back needs to overwrite the original database. If you select Yes, once the rollback is successful, the original database will be renamed RESTORE_OLD_*_OriginalDatabaseName. The new database obtained after the rollback will use the original database name.</p> <p>If you select No, once the rollback is successful, the original database remains unchanged. The name of the new database obtained after the rollback is user-defined (the default name is system-generated).</p>
Select the database to be rolled back	<p>Select the database that needs to be rolled back. The rollback of single database, multiple databases and all the databases is supported. The search function is available for quick filtering by database name. The selected databases can be renamed under the Selected Database section on the right. If they are not renamed, the created databases from the rollback will, by default, have system-generated names, with the form of prefixes and the original database names. The database name after rollback can only contain up to 128 characters, including digits, case-sensitive English letters, and special symbols (`_./()[]+=:@`). It must begin with an English letter.</p>

Scenario Two: Rollback by Backup Set

Roll back database to the specified instance

Select the target instance

Rollback to Current instance

Rollback to Existing instance

Select rollback mode

Rollback by Time Point

Rollback by Backup Set

Select replica set

mssql-2i 27:16 | 2023-12-14 05:28:06 | ▼

You can roll back a new database from a specified backup set, the range of which is determined by the data backup retention period you set.

The database name of an unarchived backup file will be displayed in the search box.

Overwrite original database ⓘ

☐ Yes

Select whether the database to be rolled back will overwrite the original database.

If you select "Yes", the original database will be renamed to "RESTORE_OLD_*_original database name" when the rollback succeeded, with the old name used for the new database.

If you don't select "Yes", the name of the original database will remain unchanged, and the new database will be named automatically by system and can be renamed later.

Select the database to be rolled back ⓘ

Select rollback database

☒ Database Name

☒ t...

Select rollback database

Da... Set Database Name After R...

te

RESTORE_2

✕

9

534_

1

✎

Hold the Shift key down to select multiple items

Save

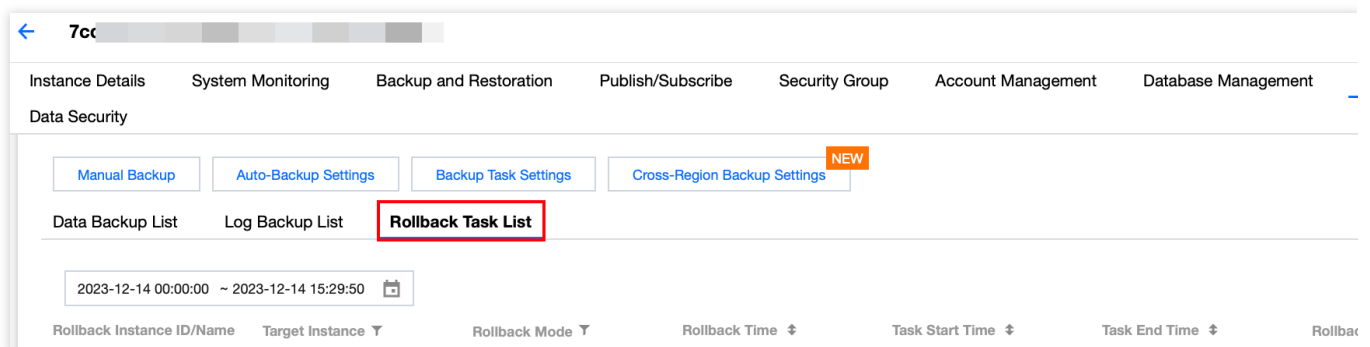
Cancel

Parameter	Description
Select the target instance	Select Rollback to Current instance .
Select rollback mode	Choose Rollback by Backup Set . You can roll back a new database from a specific backup set, with the selection range determined by the data backup retention period you have set.
Select	Select the backup set for rollback.

replica set	
Overwrite original database	<p>Select whether the database to be rolled back needs to overwrite the original database. If you select Yes, once the rollback is successful, the original database will be renamed RESTORE_OLD_*_OriginalDatabaseName. The new database obtained after the rollback will use the original database name.</p> <p>If you select No, once the rollback is successful, the original database remains unchanged. The name of the new database obtained after the rollback is user-defined (the default name is system-generated).</p>
Select the database to be rolled back	<p>Select the database that needs to be rolled back. The rollback of single database, multiple databases and all the databases is supported. The search function is available for quick filtering by database name. The selected databases can be renamed under the Selected Database section on the right. If they are not renamed, the created databases from the rollback will, by default, have system-generated names, with the form of prefixes and the original database names. The database name after rollback can only contain up to 128 characters, including digits, case-sensitive English letters, and special symbols (`_./()[]+=:@`). It must begin with an English letter.</p>

5. After confirming the rollback time or backup set and the databases that need to be rolled back, click **Save** in the pop-up window.

6. In the **Rollback Task List**, the task status changes into **In Progress**. You can view the rollback progress by clicking on the task icon in the upper right corner of the **Backup Management** page.



Rolling back to an existing instance (same region)

Last updated : 2024-01-18 17:23:30

SQL Server offers a rollback tool to perform rollback operations on instances, and reconstructs historical data through regular backups and real-time transactions. This document describes how to roll back to other instances in the same region for single database, multiple databases, or all the databases according to a time point or a backup set.

Prerequisites

To utilize the rollback feature, there must be SQL Server instances. For instructions on the creation of instances, please refer to [Creating SQL Server Instance](#).

In order to use the rollback feature, a database must exist under an instance. For instructions on the creation of a database, see [Creating a Database](#).

Points of Attention

Before initiating a rollback, please ensure there is sufficient storage space for an instance.

Ensure that no other tasks are being performed for the source instance before initiating a rollback.

The version of the target instance database after a rollback must be equal to or greater than the version of the source instance database. For instance, if the source instance is on the 2012 Enterprise version, it cannot be rolled back to a target instance on the 2008R2 Enterprise version. It can only be rolled back to an instance on the 2012 Enterprise version or above.

The architecture of the target instance after a rollback must match the architecture of the source instance. For example, a cloud disk version with two nodes can only be rolled back to a two-node instance of the same cloud disk version, not to a single-node instance of the cloud disk version or to a two-node local disk instance.

The steps are as follows:

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the **region** on the top, locate the required instance, and click **Instance ID** or **Manage** in the **Operation** column to go to the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ①
<input type="checkbox"/> mssql- ec-79d- b:00-00000000-83f	Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated	Default-VPC - Default-Subnet	Guangzhou Zone 6 Primary Replica	1433

3. Select the **Backup Management** tab, and then either click **Rollback** in the upper right corner of the backup management page or click **Rollback** in the **Data Backup List** operation column.

← 7c

Instance Details

System Monitoring

Backup and Restoration

Publish/Subscribe

Security Group

Account Management

Database Management

Data Security

Manual Backup

Auto-Backup Settings

Backup Task Settings

Cross-Region Backup Settings NEW

Data Backup List

Log Backup List

Rollback Task List

2023-12-14 00:00:00 ~ 2023-12-14 15:26:09

Separate key

Backup Start Time ①	Backup End Time ①	Backup Name	Backup Policy	Backup Mode	Upload Backup File	Retention Policy	Backup Size	Dat
2023-12-14 05:27:16	2023-12-14 05:28:06	mssql- 52716	Instance backup	Scheduled backup	Archive file	Custom	390KB	f

4. In the pop-up rollback settings interface, complete the corresponding configurations based on the selected rollback method and click **Save**.

Scenario One: Rollback by Time Point

Roll back database to the specified instance

Select the target instance

Rollback to Current instanceRollback to Existing instance

Target Instance Region

GuangzhouShanghaiBeijing

Target Instance Name

Select the target instance

Select the target instance

Select rollback mode

Rollback by Time PointRollback by Backup Set

You can roll back the database to a specific time point from a source database instance. The time range depends on the log backup retention period you set.

Set rollback time

2023-12-14 15:57:10

The rollback will be performed in Beijing time zone, with the time range available from 2023-12-07 16:13:15 to 2023-12-14 15:57:10. Please select a time point in this range.

Overwrite original database

☐ Yes

Select whether the database to be rolled back will overwrite the original database.
If you select "Yes", the original database will be renamed to "RESTORE_OLD_*_original database name" when the rollback succeeded, with the old name used for the new database.
If you don't select "Yes", the name of the original database will remain unchanged, and the new database will be named automatically by system and can be renamed later.

Select the database to be rolled back

Select rollback database

Database Name

n

Select rollback database

Da...Set Database Name After R...

RESTORE_2_7

966_7

Hold the Shift key down to select multiple items

SaveCancel

Parameter	Description
Select the target instance	Select Rollback to Existing instance .
Target Instance Region	Select the same region as the source instance.
Target	Select the target instance that you want to roll back to. You can perform a quick search based on

Instance Name	the instance ID or name. Instances across Availability Zones in the same region are supported.
Select rollback mode	Select Rollback by Time Point . You can roll back the database to a specific time point from a source database instance. The time range depends on the log backup retention period you set.
Set rollback time	Select the rollback time. A uniform rollback time can only be set for databases for the same instance.
Overwrite original database	Select whether the database to be rolled back needs to overwrite the original database. If you select Yes , once the rollback is successful, the original database will be renamed <code>RESTORE_OLD_*_OriginalDatabaseName</code> . The new database obtained after the rollback will use the original database name. If you select No , once the rollback is successful, the original database remains unchanged. The name of the new database obtained after the rollback is user-defined (the default name is system-generated).
Select the database to be rolled back	Select the database that needs to be rolled back. The rollback of single database, multiple databases and all the databases is supported. The search function is available for quick filtering by database name. The selected databases can be renamed under the Selected Database section on the right. If they are not renamed, the created databases from the rollback will, by default, have system-generated names, with the form of prefixes and the original database names. The database name after rollback can only contain up to 128 characters, including digits, case-sensitive English letters, and special symbols (<code>^_./()[]+=:@`</code>). It must begin with an English letter.

Scenario Two: Rollback by Backup Set

Roll back database to the specified instance

×

Select the target instance

Rollback to Current instance

Rollback to Existing instance

Target Instance Region

Guangzhou

Shanghai

Beijing

Target Instance Name

Select the target instance

!

Select the target instance

Select rollback mode

Rollback by Time Point

Rollback by Backup Set

Select replica set

mssql-r

|2-14 15:40:23 | 2023-12-14 15:41:17 |

▼

You can roll back a new database from a specified backup set, the range of which is determined by the data backup retention period you set.

The database name of an unarchived backup file will be displayed in the search box.

Overwrite original database ⓘ

☐ Yes

Select whether the database to be rolled back will overwrite the original database.

If you select "Yes", the original database will be renamed to "RESTORE_OLD_*_original database name" when the rollback succeeded, with the old name used for the new database.

If you don't select "Yes", the name of the original database will remain unchanged, and the new database will be named automatically by system and can be renamed later.

Select the database to be rolled back ⓘ

Select rollback database

Database Name

mje

Select rollback database

Da... Set Database Name After R...

mje RESTORE_20 1_8

↔

Hold the Shift key down to select multiple items

Save

Cancel

Parameter	Description
Select the target instance	Select Rollback to Existing instance .
Target Instance Region	Select the same region as the source instance.
Target Instance	Select the target instance that you want to roll back to. You can perform a quick search based on the instance ID or name. Instances across Availability Zones in the same region are

Name	supported.
Select rollback mode	<p>Select Rollback by Backup Set.</p> <p>Note:</p> <p>You can roll back a new database from a specific backup set, with the selection range determined by the data backup retention period you have set.</p>
Select replica set	Select the backup set for rollback.
Overwrite original database	<p>Select whether the database to be rolled back needs to overwrite the original database. If you select Yes, once the rollback is successful, the original database will be renamed <code>RESTORE_OLD_*_</code> original database name. The new database obtained after the rollback will use the original database name.</p> <p>If you select No, once the rollback is successful, the original database remains unchanged. The name of the new database obtained after the rollback is user-defined (the default name is system-generated).</p>
Select the database to be rolled back	<p>Select the database that needs to be rolled back. The rollback of single database, multiple databases and all the databases is supported. The search function is available for quick filtering by database name. The selected databases can be renamed under the Selected Database section on the right. If they are not renamed, the created databases from the rollback will, by default, have system-generated names, with the form of prefixes and the original database names.</p> <p>Note:</p> <p>The database name after rollback can only contain up to 128 characters, including digits, case-sensitive English letters, and special symbols (<code>`-_./()[]+=:@`</code>). It must begin with an English letter.</p>

Rollback to an Existing Instance (Cross-Region)

Last updated : 2024-01-18 17:23:30

SQL Server offers a rollback tool to perform rollback operations on instances, and reconstructs historical data through regular backups and real-time transactions. This document describes how to roll back to other instances across different regions for single database, multiple databases, or all the databases according to a time point or a backup set.

Note:

Rolling back to another instance across regions refers to realizing cross-region recovery through cross-region backups for the purpose of remote disaster recovery. For example, if an instance in region A has enabled the cross-region backup function and its backup files are stored in region B, when the instance encounters a failure or damage, the cross-region backup files can be restored to an instance in region B by using the rollback function.

Prerequisites

To utilize the rollback feature, there must be SQL Server instances. For instructions on the creation of instances, please refer to [Creating SQL Server Instance](#).

In order to use the rollback feature, a database must exist under an instance. For instructions on the creation of a database, see [Creating a Database](#).

Cross-region backups must have been enabled. Cross-region backup files must have been generated. Please refer to [Cross-Region Backup](#).

Points of Attention

Before initiating a rollback, please ensure there is sufficient storage space for an instance.

Ensure that no other tasks are being performed for the source instance before initiating a rollback.

The version of the target instance database after a rollback must be equal to or greater than the version of the source instance database. For instance, if the source instance is on the 2012 Enterprise version, it cannot be rolled back to a target instance on the 2008R2 Enterprise version. It can only be rolled back to an instance on the 2012 Enterprise version or above.

The architecture of the target instance after a rollback must match the architecture of the source instance. For example, a cloud disk version with two nodes can only be rolled back to a two-node instance of the same cloud disk version, not to a single-node instance of the cloud disk version or to a two-node local disk instance.

The steps are as follows:

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the **region** on the top, locate the required instance, and click **Instance ID** or **Manage** in the **Operation** column to go to the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ
<input type="checkbox"/> mssql- ec-79d- b:83f	Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated	Default-VPC - Default-Subnet	Guangzhou Zone 6 Primary Replica	1433

3. Select the **Backup Management** tab, and then either click **Rollback** in the upper right corner of the backup management page or click **Rollback** in the **Data Backup List** operation column.

← 7c

Instance DetailsSystem MonitoringBackup and RestorationPublish/SubscribeSecurity GroupAccount ManagementDatabase Management

Data Security

Manual BackupAuto-Backup SettingsBackup Task SettingsCross-Region Backup SettingsNEW

Data Backup ListLog Backup ListRollback Task List

2023-12-14 00:00:00 ~ 2023-12-14 15:26:09 Separate key

Backup Start Time ⓘ	Backup End Time ⓘ	Backup Name	Backup Policy	Backup Mode	Upload Backup File	Retention Policy	Backup Size	Da
2023-12-14 05:27:16	2023-12-14 05:28:06	mssql-52716	Instance backup	Scheduled backup	Archive file	Custom	390KB	1

4. In the pop-up rollback settings interface, complete the corresponding configurations based on the selected rollback method and click **Save**.

Scenario One: Rollback by Time Point

Roll back database to the specified instance

Select the target instance

Rollback to Current instance

Rollback to Existing instance

Target Instance Region

Guangzhou

Shanghai

Beijing

Target Instance Name

Select the target instance

Select the target instance

Select rollback mode

Rollback by Time Point

Rollback by Backup Set

Set rollback time

2023-12-14 15:57:10

The rollback will be performed in Beijing time zone, with the time range available from 2023-12-07 16:16:12 to 2023-12-14 15:57:10. Please select a time point in this range.

Overwrite original database

☐ Yes

Select whether the database to be rolled back will overwrite the original database.
If you select "Yes", the original database will be renamed to "RESTORE_OLD_"_original database name" when the rollback succeeded, with the old name used for the new database.
If you don't select "Yes", the name of the original database will remain unchanged, and the new database will be named automatically by system and can be renamed later.

Select the database to be rolled back

Select rollback database

Database Name

m

Select rollback database

Da... Set Database Name After R...

RESTORE_2I 96f

Hold the Shift key down to select multiple items

Save

Cancel

Parameter	Description
Select the target instance	Select Rollback to Existing instance .
Target Instance Region	Select rollback to the region where other instances are.
Target Instance Name	Select the cross-region target instance after a rollback. You can quickly search by instance ID or instance name. Instances across Availability Zones in the selected region are supported.
Select	Select Rollback by Time Point .

©2013-2022 Tencent Cloud. All rights reserved.

Page 327 of 446

rollback mode	You can roll back the database to a specific time point from a source database instance. The time range depends on the log backup retention period you set.
Set rollback time	Select the rollback time. A uniform rollback time can only be set for databases for the same instance.
Overwrite original database	Select whether the database to be rolled back needs to overwrite the original database. If you select Yes , once the rollback is successful, the original database will be renamed RESTORE_OLD_*_OriginalDatabaseName . The new database obtained after the rollback will use the original database name. If you select No , once the rollback is successful, the original database remains unchanged. The name of the new database obtained after the rollback is user-defined (the default name is system-generated).
Select the database to be rolled back	Select the database that needs to be rolled back. The rollback of single database, multiple databases and all the databases is supported. The search function is available for quick filtering by database name. The selected databases can be renamed under the Selected Database section on the right. If they are not renamed, the created databases from the rollback will, by default, have system-generated names, with the form of prefixes and the original database names. The database name after rollback can only contain up to 128 characters, including digits, case-sensitive English letters, and special symbols (`-_./()[]()+=::@`). It must begin with an English letter.

Scenario Two: Rollback by Backup Set

Roll back database to the specified instance

Select the target instance

Rollback to Current instance

Rollback to Existing instance

Target Instance Region

Guangzhou

Shanghai

Beijing

Target Instance Name

Select the target instance

Select the target instance

Select rollback mode

Rollback by Time Point

Rollback by Backup Set

Select replica set

mssql-r

2023-12-14 15:41:17

Overwrite original database

Yes

Select whether the database to be rolled back will overwrite the original database.

If you select "Yes", the original database will be renamed to "RESTORE_OLD_*_original database name" when the rollback succeeded, with the old name used for the new database.

If you don't select "Yes", the name of the original database will remain unchanged, and the new database will be named automatically by system and can be renamed later.

Select the database to be rolled back

Select rollback database

Select rollback database

Save

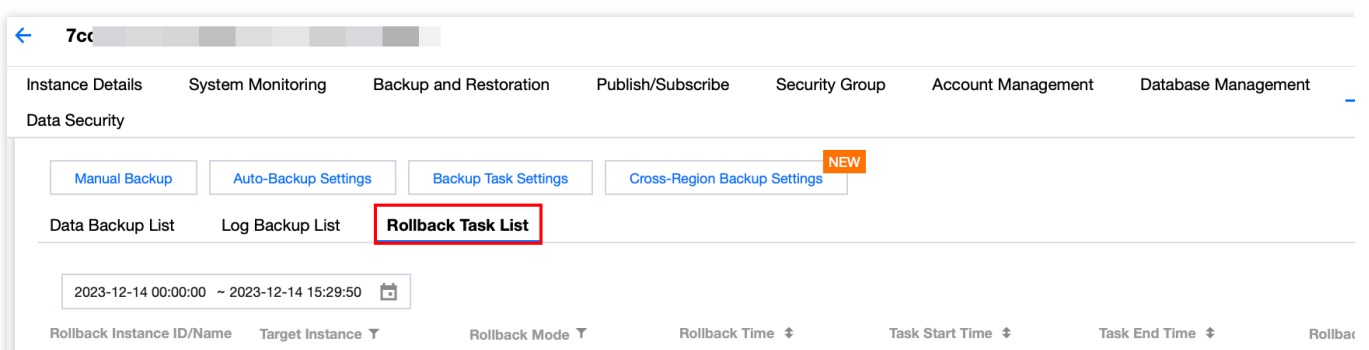
Cancel

Parameter	Description
Select the target instance	Select Rollback to Existing instance .
Target Instance Region	Select rollback to the region where other instances are.
Target Instance Name	Select the cross-region target instance after a rollback. You can quickly search by instance ID or instance name. Instances across Availability Zones in the selected region are supported.

Select rollback mode	<p>Choose Rollback by Backup Set.</p> <p>Note:</p> <p>You can roll back a new database from a specific backup set, with the selection range determined by the data backup retention period you have set.</p>
Select replica set	Select the backup set for rollback.
Overwrite original database	<p>Select whether the database to be rolled back needs to overwrite the original database. If you select Yes, once the rollback is successful, the original database will be renamed <code>RESTORE_OLD_*_OriginalDatabaseName</code>. The new database obtained after the rollback will use the original database name.</p> <p>If you select No, once the rollback is successful, the original database remains unchanged. The name of the new database obtained after the rollback is user-defined (the default name is system-generated).</p>
Select the database to be rolled back	<p>Select the database that needs to be rolled back. The rollback of single database, multiple databases and all the databases is supported. The search function is available for quick filtering by database name. The selected databases can be renamed under the Selected Database section on the right. If they are not renamed, the created databases from the rollback will, by default, have system-generated names, with the form of prefixes and the original database names.</p> <p>Note:</p> <p>The database name after rollback can only contain up to 128 characters, including digits, case-sensitive English letters, and special symbols (<code>`-_./()[]()+=:@`</code>). It must begin with an English letter.</p>

5. After confirming the rollback time or backup set and the databases that need to be rolled back, click **Save** in the pop-up window.

6. In the **Rollback Task List**, when the task status changes into **In Progress**, you can view the rollback progress by clicking on the task icon at the top right corner of the **Backup Management** page.



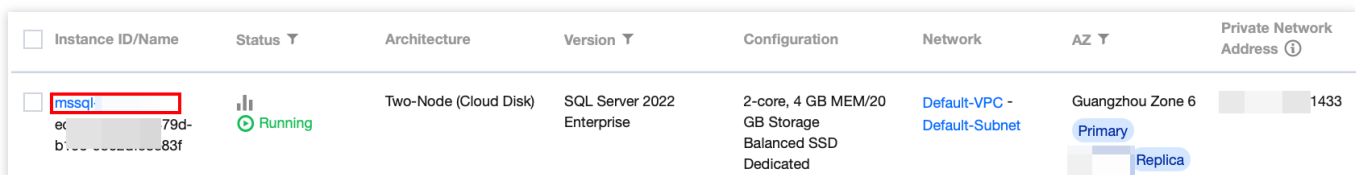
Viewing Rollback Task List

Last updated : 2024-01-18 17:23:30

You can view the rollback task progress for instances or view the details of historical rollback tasks through the **Rollback Task List**. This document provides instructions on how to view the **Rollback Task List** via the console.

Viewing Rollback Task List

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the **region** on the top, locate the required instance, and click **Instance ID** or **Manage** in the **Operation** column to go to the instance management page.



<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ
<input type="checkbox"/> mssql ec-79d- b:83f	Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated	Default-VPC - Default-Subnet	Guangzhou Zone 6 Primary Replica	1433

3. Click **Backup Management** and then select the **Rollback Task List** on the Backup Management page.

<

4. The fields you can view are as follows:

Field	Description
Rollback Instance ID/Name	Source Instance ID/Name.
Target Instance	Filters according to the current instances, existing instances, or completely new instances.
Rollback Mode	Filters according to Rollback by Time Point or Rollback by Backup Set .
Rollback Time	Supports display in ascending or descending order.
Task Start Time	Supports display in ascending or descending order.
Task End time	Supports display in ascending or descending order.

Rollback Region	Displays the region to be rolled back to.
Status	Displays the rollback task status.
Operation	Supports the deletion of rollback tasks.

Log Management

Querying and Downloading Blocking and Deadlock Events

Last updated : 2024-01-18 17:23:30

TencentDB for SQL Server allows you to record blocking and deadlock events and download corresponding log files. This helps you locate and optimize SQL statements that cause blocking and deadlock.

This document describes how to query and download blocking and deadlock events.

Note:

Blocking and deadlock events are supported only in 2012, 2014, 2016, 2017, and 2019 Enterprise Editions, which are not available in 2008 R2 Enterprise Edition.

The log is stored in the Beijing time zone (UTC+8) by default. If the default time zone is modified, it will be stored in the time zone of the instance server, but its storage time will still be in Beijing time in the console.

Prerequisites

To query and download blocking and deadlock events, you need to enable the collection of blocking events and deadlock events first.

Background

To ensure data consistency in a database, a resource is not released while it is still being modified, so that it can't be accessed or modified by other concurrent sessions. However, it may be occupied for an extended period of time if there are slow SQL queries or other exceptions, resulting in a blocking event where it can't be accessed by other sessions.

On the other hand, if multiple transactions compete for resources, for example, when two transactions hold resources without releasing them and attempt to access the resource held by each other, a deadlock event will occur.

To address these issues, TencentDB for SQL Server provides a feature to record the blocking and deadlock events in the console. Once this feature is enabled, you can quickly identify blocking and deadlock events occurring in the database, helping you locate and optimize the execution SQL statements that caused the problem.

Use Limits

By default, the collection of blocking and deadlock events is disabled, but it can be enabled manually. Once enabled, such events will be logged in the log files, along with the details of the SQL execution that caused the blocking and deadlock.

After the collection of blocking and deadlock events is enabled, the collection threshold is set to 1000 ms (or 1s) by default, which can be set to anywhere between 1000 and 86,400,000 ms. A SQL running beyond the threshold will be logged as a blocking SQL and a deadlock SQL.

By default, blocking and deadlock events are collected every 5 minutes. This means that any SQL query that exceeds 1s (the default threshold) will be recorded within a 5-minute interval.

Blocking and deadlock events are retained for 7 days by default and automatically deleted upon expiration.

Enabling the collection of blocking and deadlock events

Note:

The collection of blocking events and deadlock events can only be enabled or disabled simultaneously. It is currently not possible to configure the collection settings separately for each event type.

Option 1. Enable through operation log settings

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the instance ID or **Manage** in the **Operation** column to enter the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ	E
<input type="checkbox"/> mssql-53b5-4d-3e	Running	Single-Node (Cloud Disk)	SQL Server R2 Enterprise	2-core, 4 GB MEM/20 GB Storage Premium Cloud Disk Dedicated		Beijing Zone 4	1:1433	F

3. On the **Instance Management** page, select the **Operation log** tab.
4. Click **Operation Log Settings**.

Operation Log	Parameter Configuration	Data Security
<div>Operation Log Settings</div> <div> Slow Query Log Blocking Event Deadlock Event Primary/Replica Switch Log </div> <div> 2023-11-06 00:00:00 ~ 2023-11-06 17:31:49 </div> <div> File Name Start Time of File Generation ⓘ End Time of File Generation ⓘ File Size </div>		

5. In the pop-up window, complete the following configuration and click **Save**.

Operation Log Settings

Slow Query Log Settings

Collection ⓘ

☒ Enable

Collection Threshold ⓘ

1000

ms

Collection Frequency ⓘ

Every 5 minutes

Retention Period

7 days

Blocking and Deadlock Event Settings

Collection ⓘ

☒ Enable

Collection Threshold ⓘ

1000

ms

✓

Range: [1,000-86,400,000] ms

Collection Frequency ⓘ

Every 5 minutes

Retention Period

7 days

Save

Cancel

Parameter	Description
Collection	Toggle on this fswitch.
Collection Threshold	Set a collection threshold to anywhere between 1000 and 86,400,000 ms.
Collection Frequency	It is every 5 minutes by default and can't be modified.
Retention Period	It is 7 days by default. The events will be automatically deleted upon expiration.

Option 2. Enable by setting the parameter blocked process threshold

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID to access the instance management page.
2. On the **Instance Management** page, select **Parameter Configuration** > **Parameter Settings** tab. Then, find the parameter `blocked process threshold` , and click



to set a non-zero value on its **Current Value** column.

Note:

The default value for the parameter `blocked process threshold` is 0, which means that blocking events and deadlock events are not collected.

The `blocked process threshold` parameter can be set to a value between 0 and 86,400s.


If the current running value of the parameter `blocked process threshold` is not 0, it indicates that blocking and deadlock events are being collected. The collection switch is toggled on in the [Operation Log Settings] correspondingly.

Querying and downloading blocking and deadlock events

1. On the **Instance Management** page, select **Operation Log** > **Blocking Event**, or **Operation Log** > **Deadlock Event** to view the corresponding event lists.

You can view the following fields: **File Name**, **Start Time of File Generation**, **End Time of File Generation**, **File Size**, and **Operation (Download)**.

You can search for slow logs generated in the last 5/15/30 minutes, last 1/3/24 hours, today, yesterday, last 3/7/30 days, or a custom time range.



2. Click **Download** in the **Operation** column to download the blocking events or deadlock events file.

Querying and Downloading Slow Query Log

Last updated : 2024-01-18 17:23:30

The slow query log records the query statements that take more time than the specified value to execute in TencentDB for SQL Server. It enables you to find out inefficient query statements for optimization. Typically, it is a SQL statement for troubleshooting and an important feature for checking the performance of the current TencentDB for SQL Server instance.

This document describes how to query and download a slow query log in the console.

Note:

The log is stored in the Beijing time zone (UTC+8) by default. If the default time zone is modified, it will be stored in the time zone of the instance server, but its storage time will still be in Beijing time in the console.

Use Limits

The collection of slow query is enabled by default and cannot be disabled.



The collection threshold of slow query is 1s (or 1,000 ms) by default, and SQL executions exceeding 1s will be recorded as a slow query log.

The slow query logs are collected every 5 minutes by default, and SQL statements exceeding 1s will be recorded.

The slow query logs are retained for 7 days by default and will be automatically deleted upon expiration.

Directions

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the region at the top, find the target instance, and click the instance ID or **Manage** in the **Operation** column to enter the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Project ▾	Architecture	Version ▾	Configuration	Network	AZ ▾
<input type="checkbox"/> mssql-53328b66-05e6-444d-b5d2	 Running	Default Project	Single-Node (Cloud Disk)	SQL Server 2008 R2 Enterprise	2-core, 4 GB MEM/20 GB Storage High-Performance Cloud Disk Dedicated		Beijing Zone 4


3. On the instance management page, select the **Operation Log** tab > **Slow Query Log** to see slow log list.

You can view the following fields: **File Name**, **Start Time of File Generation**, **End Time of File Generation**, **File Size**, and **Operation (Download)**.

You can search for slow logs generated in the last 5/15/30 minutes, last 1/3/24 hours, today, yesterday, last 3/7/30 days, or a custom time range.

[Instance Details](#)[System Monitoring](#)[Backup and Restoration](#)[Security Group](#)[Account Management](#)[Database Management](#)[Parameter Configuration](#)[Data Security](#)

Slow Query Log

2023-07-07 00:00:00 ~ 2023-07-07 10:05:36 

File Name	Start Time of File Generation ⓘ	End Time of File Generation ⓘ	File Size
-----------	---------------------------------	-------------------------------	-----------

4. Click **Download** in the **Operation** column to download slow query log files.

Querying primary and standby switch logs

Last updated : 2024-01-18 17:23:30

SQL Server supports the records for primary and secondary switch logs of an instance. Through that, you can understand the specific time and switch methods of the primary and secondary switch to verify the success of the switch, and optimize switch policies.

This document describes how to query primary and secondary switch logs.

The steps are as follows:

1. Log in to the [TencentDB for SQL Server console](#).
2. Select the **region** on the top, locate the instance that needs to query the primary and secondary switch logs, and click **Instance ID** or **Manage** in the **Operation** column to go to the instance management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ
<input type="checkbox"/> mssql- ec-79d- b.....83f	Running	Two-Node (Cloud Disk)	SQL Server 2022 Enterprise	2-core, 4 GB MEM/20 GB Storage Balanced SSD Dedicated	Default-VPC - Default-Subnet	Guangzhou Zone 6 Primary Replica	1433

3. On the instance management page, choose **Operation Log > Primary/Replica Switch Log** to view the list of primary and secondary switch logs.

You can view the following fields: **Switch Event ID**, **Switch Start Time**, **Switch End Time**, and **Switch Mode** (automatic switch by system, manual switch).

You can filter logs by using a time-based search.

Instance Details	System Monitoring	Backup and Restoration	Publish/Subscribe	Security Group	Account Management	Database Management
Parameter Configuration	Data Security					
Operation Log Settings						
Slow Query Log	Blocking Event	Deadlock Event	Primary/Replica Switch Log			
2023-12-14 00:00:00 ~ 2023-12-14 16:21:56						
Switch Event ID		Switch Start Time		Switch End Time		

Manual Creation of the Latest Blocking and Deadlock Events

Last updated : 2024-07-31 09:44:54

TencentDB for SQL Server supports manually creating the latest blocking and deadlock events, provided the collection of blocking and deadlock events is enabled. You can immediately create and download the blocking and deadlock event files when needed. This document introduces how to manually create the latest blocking and deadlock events through the console.

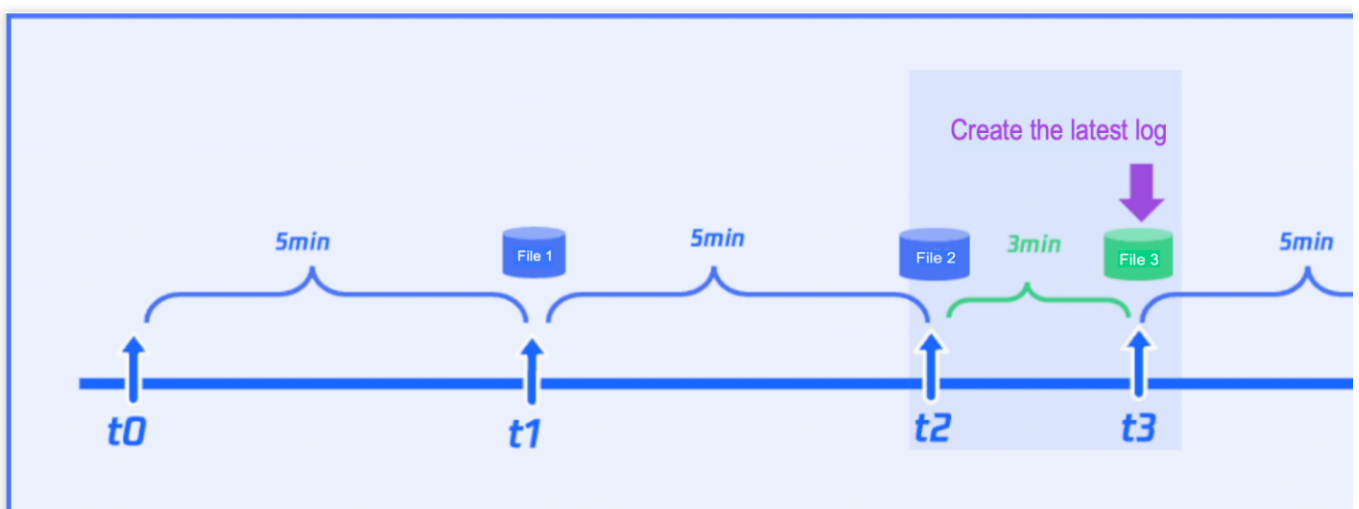
Note:

Only Enterprise editions 2012, 2014, 2016, 2017, 2019, and 2022 support manually creating the latest blocking and deadlock events. The 2008 R2 Enterprise edition does not support this feature.

The log time is set to Beijing Time (UTC+8) by default. If the default timezone is modified, the log data storage time will be displayed in the instance server's timezone, but the console will show Beijing Time (UTC+8).

Operation Scenarios

After the collection of blocking and deadlock events is enabled for TencentDB for SQL Server, the blocking and deadlock event files are collected every 5 minutes by default. That is, within every 5-minute interval, all SQL executions exceeding the collection threshold will be recorded and a file generated. You can also manually generate and download the blocking and deadlock event files immediately as needed. For example, within a log file generation cycle (every 5 minutes), if the latest log has been collected for 3 minutes, you can immediately download the generated log files for these 3 minutes using the manual creation feature.



Prerequisites

Enabling the Collection of Blocking Events and Deadlock Events.

Manually Creating the Latest Blocking and Deadlock Events

1. Log in to the [SQL Server console](#).
2. Select the region at the top, find the target instance, click **Instance ID** or click **Manage** in its **Operation** column to enter the Instance Management page.

<input type="checkbox"/> Instance ID/Name	Status ▾	Architecture	Version ▾	Configuration	Network	AZ ▾	Private Network Address ⓘ	Billing Mod
<input type="checkbox"/> mssql-534d-b5e	Running	Single-Node (Cloud Disk)	SQL Server R2 Enterprise	2-core, 4 GB MEM/20 GB Storage Premium Cloud Disk Dedicated		Beijing Zone 4	1:143	Pay as You (

3. On the Instance Management page, select the **Operation Log** tab.
4. Click **Create Latest Log**.

Backup Management	Operation Log	Parameter Configuration	Data Security
Operation Log Settings	Create Latest Log		
Slow Query Log	Blocking Event	Deadlock Event	

5. In the pop-up window, click **Save**.

Note:

Creating the latest log will immediately cut off and upload the currently generating file that has not yet met the packaging time. For example: if the collection frequency is set to generate one log file every 5 minutes, and the latest file has only been collected for 3 minutes, clicking **Create Latest Log** will immediately cut off, package, and upload the current 3-minute log for the user to download. Subsequent log files will continue to be generated from the cut-off time according to the predetermined collection frequency.

Parameter	Description
Log type	The default values are blocking events and deadlock events. Currently, only simultaneous creation is supported; creating blocking events or deadlock events individually is not supported.
Collection threshold	SQL executions beyond the collection threshold will be recorded as Blocking SQL and Deadlock SQL. By default, the threshold here is the same as the settings when enabling blocking and deadlock events.
Retention Period	The default is 7 days.

Subsequent Operations

[Querying and Downloading Blocking Events and Deadlock Events](#)

FAQs

Why is the feature option for creating the latest log not clickable in the console?

1. Your instance may not have collection of blocking and deadlock events enabled. You need to enable event collection first. Please refer to [Enabling Collection of Blocking and Deadlock Events](#).
2. Your instance version may be 2008 R2 Enterprise edition. Currently, only 2012, 2014, 2016, 2017, 2019, 2022 Enterprise editions support manual creation of the latest blocking and deadlock logs.

How to modify the collection threshold?

Due to the SQL collection threshold recorded in the manually created latest log files being related to the settings when enabling the collection of blocking and deadlock events, you can modify it in **Operation Log > Operation Log Settings**.

Publish-Subscribe

Overview of Publish-Subscribe

Last updated : 2024-01-18 17:23:30

Overview

TencentDB for SQL Server supports the native **Publish/Subscribe** replication function of Microsoft SQL Server. Users can create, modify, and delete publishing and subscription servers on the [TencentDB for SQL Server console](#), meeting the business requirements for data replication and synchronization.

Note :

SQL Server single-node (formerly the basic version) does not support the **Publish/Subscribe** feature.

Concepts

SQL Server employs terminology from the publishing industry to represent elements of a replication topology, which includes the publishing server, distributing server, subscribing server, publications, projects, and subscriptions. The Microsoft SQL Server replication can be understood based on the perception of the concepts of magazines:

The magazine publisher (Publishing Server) creates one or more publications (Publication).

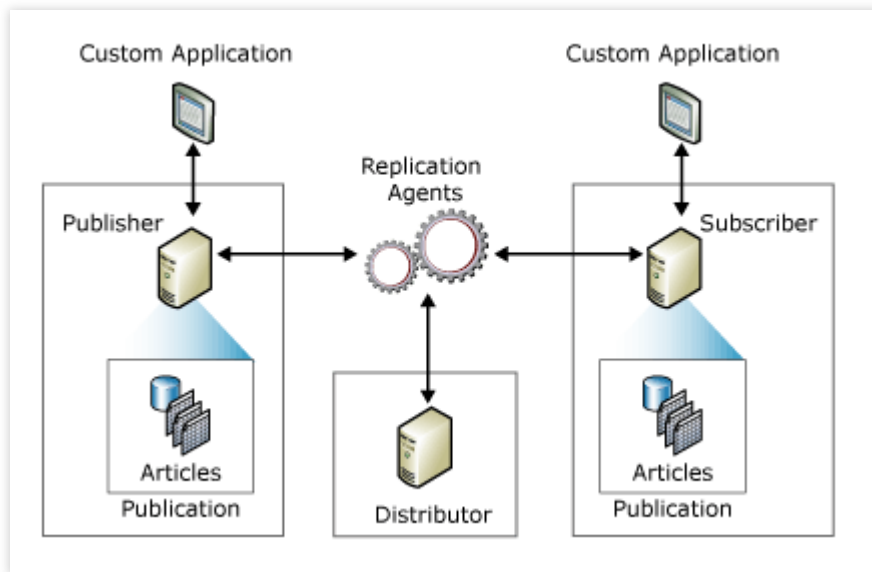
The publications (Publication) comprise articles (Project).

The publisher (Publishing Server) may directly distribute the magazines or resort to a distributor (Distributing Server).

The subscriber (Subscribing Server) receives the subscribed publications (Publication).

Architecture

The replication topology defines the relationships between servers and data replicas, and illustrates the logic that determines how data flows between servers. Various replication processes, termed as agents, are devoted to copying and moving data between publishing and subscribing servers. The components and processes involved in replication are as follows:.



Publishing Server

A Publishing Server is a type of database instance that provides data to other locations via replication. A Publishing Server can have one or multiple publications, each defining a set of objects and data to be replicated with logical relationships.

Distributing Server

A Distributing Server is another type of database instance, functioning as a storage hub. It is used for the replication of the specific data associated with one or multiple publishing servers. Each Publishing Server is associated with a single database (distributing database) in the Distributing Server.

The distributing database can store replication state data and metadata regarding publication. In certain scenarios, it queues the data moving from the Publishing Server to the Subscribing Server.

In many scenarios, a single database server instance serves as the two roles of both the Publishing Server and the Distributing Server, which is called Local Distributing Server. When the Publishing Server and Distributing Server are configured according to their respective database server instances, the Distributing Server is named Remote Distributing Server.

Subscribing Server

The Subscription Server is a database instance that receives replicated data. This server can receive the published data from multiple Publishing Servers. Depending on the selected replication type, the Subscribing Server can either transmit data alterations back to the Publishing Server or republish data to other Subscribing Servers.

Project

A project identifies the database objects included in a publication. A single publication can encompass various types of projects.

Publication

Publication is a collection of one or more projects within a database. Grouping multiple projects into a publication facilitates the designation of a set of database objects and data to be replicated as a unit with logical relationships.

Subscription

A subscription is a request to deliver a publication replica to a Subscribing Server. It defines the publication to be received and the time as well as the location to be received.

Feature Overview

1. The SQL Server Publish/Subscribe function typically utilizes transaction replication by default.

Transaction replication usually begins with a snapshot of the published database objects and data. After the initial snapshot is created, data changes and schema modifications made on the Publishing Server are usually delivered to the Subscribing Server as the changes occur, providing an almost real-time replication service. Data modifications are applied to the Subscribing Server in the order they occurred on the Publishing Server and within their transaction boundaries, thereby ensuring transactional consistency within the publication.

Transaction replication is typically used in server-to-server environments and proves appropriate in the following scenarios:

When there is an incremental change, it can be delivered to the Subscribing Server.

The application requires a minimal latency period between changes occurring on the Publishing Server and those changes arriving at the Subscribing Server.

The application needs access to intermediate data states. For instance, if a row changes five times, transaction replication would allow the application to respond to each change (such as invoking triggers) rather than simply responding to the final data state of the row.

There's a high volume of insertion, update, and deletion activities in the Publishing Server.

2. SQL Server Publish/Subscribe configuration uses a remote Distributing Server to ensure that after a failure switch of the Publishing Server, the **Publish/Subscribe** link automatically recovers, making the process imperceptible to the Subscribing Server. However, after a failure of the Subscribing Server, the link requires manual rectification.


Feature Limits

This feature is available only when the publishing instances and subscribing instances are TencentDB for SQL Server instances.

A read-only instance cannot serve as a publishing or subscribing server.

The publishing and subscribing instances must have the same version and situate in the same region, though they can be in different availability zones.

Data tables without a primary key cannot be subscribed to. The following code can be employed to examine whether the database to be published contains tables not equipped with a primary key.



```
use dbname
select name from sys.sysobjects where xtype='U' and id not in(select parent_obj fro
```

When a database with the same name exists in both the publishing and subscribing instances, such a database cannot be subscribed to.

Once a **Publish/Subscribe** link is established, if a database in the link is deleted, the established **Publish/Subscribe** link will accordingly be deleted.

If either the publishing or subscribing instance is terminated, the corresponding **Publish/Subscribe** link will also be deleted.

Managing Publish-Subscribe

Last updated : 2024-07-30 16:06:03

Operation Scenarios

You can manage **Publish/Subscribe** via the SQL Server console, including creating, viewing, and deleting **Publish/Subscribe** tasks.

Note :

During the enabling period of the publish-subscribe feature, it is advisable to refrain from manipulating the data in the subscription database, as this may lead to inconsistencies in data synchronization.

Feature Limits


This feature is available only when the publishing instances and subscribing instances are TencentDB for SQL Server instances.

Only database-level publication and subscription are supported.

A read-only instance cannot serve as a publishing or subscribing server.

The publishing and subscribing instances must have the same version and situate in the same region, though they can be in different availability zones.

Data tables without a primary key cannot be subscribed to. The following code can be employed to examine whether the database to be published contains tables not equipped with a primary key.



```
use dbname
select name from sys.sysobjects where xtype='U' and id not in(select parent_obj fro
```

When a database with the same name exists in both the publishing and subscribing instances, such a database cannot be subscribed to.

Once a **Publish/Subscribe** link is established, if a database in the link is deleted, the established **Publish/Subscribe** link will accordingly be deleted.

If either the publishing or subscribing instance is terminated, the corresponding **Publish/Subscribe** link will also be deleted.


Prerequisites

The SQL Server instances have been established. Ensure a minimum of two instances. For more information, refer to [Creating SQL Server Instance](#).

The steps are as follows:

Establishing Publish-Subscribe Relationships

1. Log into the [SQL Server console](#). In the instance list, click the **Instance ID** or **Manage** in the **Operation** column to access the instance management page.
2. On the instance management page, select the **Publish/Subscribe** tab and then click **Create**.



3. On the **Set publishing and subscription instances** page, enter the **Publish/Subscribe Name**, select the **Subscribe Instance ID**, and click **Next**.

Note:

A read-only instance cannot serve as a publishing or subscribing server.

The publishing and subscribing instances must have the same version and situate in the same region, though they can be in different availability zones.

[←](#) **Create pub/sub project**

1 Set publishing and subscription instances

>

2 Select pub/sub project

>

3 Confirm pub/sub project

Basic Info

Pub/Sub Name *

Publishing Settings

Version *

Region *

Publish Instance ID *

Subscription Settings

Version *

Region *

Subscribe Instance ID * 31 ▼

4. On the **Select pub/sub project** page, **select a pub/sub database** (you can select multiple databases), then click **Go Next and Confirm Configuration**.

Note:

You can configure up to 80 databases to be published/subscribed to in each **Publish/Subscribe** operation by default.

1 Set publishing and subscription instances > 2 Select pub/sub project > 3 Confirm pub/sub project

Select a pub/sub database

Database
<input checked="" type="checkbox"/> Database
<input checked="" type="checkbox"/> n

Hold the Shift key down to select multiple items

Selected (1)

Database
n

Previous Go Next and Confirm Configuration

5. On the **Confirm pub/sub project** page, confirm the configuration information for **Publish/Subscribe**, then click **Complete and Start Configuration**.

1 Set publishing and subscription instances > 2 Select pub/sub project > 3 Confirm pub/sub project

Add Subscription

Publish Database	Subscribe Database
n	m

Previous Complete and Start Configuration

You can view the progress of the **Pub/Sub Creation** via **Running Tasks** in the upper right corner of the **Publish/Subscribe** page.

Head-Only Instance

Authorization Instructions

Add Batch D

Terminate/Return Rollback

Running Tasks Show All Tasks

Task ID	Task Name	Status	Progress	Start time
ecc3	Pub/Sub Creation	Running	20%	2023-12-14 16:45:26

Viewing Publish/Subscribe Status

1. Log in to the [TencentDB for SQL Server console](#), on the instance list, click the **Instance ID** of the publishing instances or the subscribing instances to access the instance management page.
2. On the instance management page, select the **Publish/Subscribe** tab to view the created **Publish/Subscribe** relationship. The key information includes the basic information of the publishing/subscribing instances, the statuses, and the last synchronization time.

Publish name	Publishing instance ID	Publish instance name	Publish instance IP	Subscription instance ID	Subscribe instance name	Subscribe instance IP	Status	Last Synced Time
<input type="checkbox"/> Publish name: test	mssql-r...	ecc3e...	...	mssql...	Initializing	2023-12-...

Deleting Publish/Subscribe Relationships

1. Log in to the [TencentDB for SQL Server console](#), on the instance list, click the **Instance ID** of the publishing instances or the subscribing instances to access the instance management page.
2. On the instance management page, select the **Publish/Subscribe** tab, select the required publishing/subscribing instances, and click **Delete**. You can also delete the instances in batches.

Note:

Once a **Publish/Subscribe** link is established, if a database in the link is deleted, the established **Publish/Subscribe** link will accordingly be deleted.

If either the publishing or subscribing instance is terminated, the corresponding **Publish/Subscribe** link will also be deleted.

SSIS

Overview

Last updated : 2024-01-18 17:23:30

Overview

SSIS, SSAS, and SSRS are three key components for SQL Server to implement business intelligence (BI).

SQL Server Integration Services (SSIS) provides enterprise-grade data integration and conversion solutions to extract, transform, and load (ETL) data from various sources.

SQL Server Analysis Services (SSAS) is a data analysis engine used to create cubes (multidimensional data models).

SQL Server Reporting Services (SSRS) is a report tool used to create, deploy, and manage mobile and paginated reports.

TencentDB for SQL Server has released the business intelligence server feature, which provides a complete set of BI solutions integrating data storage, ETL, and visual analysis and supports SSIS. SSIS can be used to sustain complex business scenarios, such as merging data from heterogeneous data stores, cleansing and standardizing data, populating data warehouses and datasets, transforming data for complex business logic, supporting management features, and automating data loading. It helps meet your diversified needs in various use cases, including BI analysis, high-value data mining, and primary data management system setup.

SSIS can extract and transform data from various data sources and load the data into one or multiple targets. SSIS capabilities in Tencent Cloud currently can be used for TencentDB for SQL Server and flat files (with .txt, .csv, .xlsx, and .xls extensions).

To use the SSIS capabilities in TencentDB for SQL Server, you need to use the SSIS engine in a business intelligence server to deploy a project first.

Use Cases

Merging data from heterogeneous data stores

Data is usually stored in many different data storage systems, and you often need to extract data from these sources and merge it into a single consistent dataset. This process faces various problems, including complex and different traditional systems, data formats, and preprocessing flows.

SSIS can use .NET and OLE DB access APIs to connect to TencentDB and use an ODBC driver to connect to multiple legacy databases. It can also connect to flat files and Excel files. In addition, it has some source components to extract data from different data sources and provides the data transformation feature. After transforming data into

compatible formats, it can further merge data physically into one target database. Then, it can load the data into flat files and SQL Server databases.

Cleansing and standardizing data

Various data sources use different conventions and standards, and different business processing jobs need to be executed during data loading. Therefore, whether data is loaded into an online transaction processing (OLTP) or online analytic processing (OLAP) database, an Excel spreadsheet, or a file, it needs to be cleaned and standardized before it is loaded.

SSIS includes built-in transformations that you can add to packages to clean and standardize data, change the case of data, convert data to a different type or format, or create new column values based on expressions. For example, a package can concatenate first and last name columns into a single full name column, and then change the characters to uppercase. A package can also clean data by replacing the values in columns with values from a reference table, using either an exact lookup or fuzzy lookup to locate values in the reference table.

Processing complex business logic during data transformation

A data transformation process requires built-in logic to respond dynamically to the data it accesses and processes. The data may need to be summarized, converted, and distributed based on data values. The process may even need to reject data based on an assessment of column values. SSIS offers many types of relevant jobs:

Merging data from multiple data sources.

Evaluating data and applying data conversions.

Splitting a dataset into multiple datasets based on data values.

Applying different aggregations to different subsets of a dataset.

Loading subsets of the data into different or multiple destinations.

Automating administrative features and data loading

SSIS provides components to automate management, such as copying TencentDB for SQL Server databases and their contained objects, copying TencentDB for SQL Server objects, loading data, and setting the scheduling cycle and frequency of SSIS tasks.

Benefits

High-value data mining

SSIS aggregates discrete and partially structured data in an enterprise at different standards by cleansing and processing the data. This helps tap into the value of data and form an enterprise-level unified database, which serves as a high-quality data source for enterprise decision making.

Primary data management system setup

Legacy business systems may have various problems such as complex and different data sources, formats, and preprocessing flows required for merging. SSIS can extract, transform, load, and merge data from multiple storage systems into one target database for unified primary business database setup. This not only facilitates internal data maintenance and management but also reduces the storage and maintenance costs otherwise incurred by many different databases.

BI analysis

SSIS provides a complete set of BI analysis solutions ranging from data storage and ETL to visual analysis. With SSIS, you can perform such operations on different data sources at one stop and then easily implement real-time self-service visual data analysis throughout the entire linkage. SSIS also helps you set up an enterprise data middleend to guide refined data-driven business operations.

Data integration for internal business systems

SSIS greatly improves your efficiency, accuracy and system performance of data collection and cleansing while simplifying the entire data aggregation and analysis process. You can configure SSIS to automatically schedule data ETL jobs, which facilitates data integration with your internal business systems.

Use Limits

Currently, the business intelligence server feature is in beta test and can be used free of charge. During the beta test, you can purchase only one business intelligence server with the 2-core 4 GB MEM specification in each region and up to three under each root account. Billing will start in pay-as-you-go mode after the beta test ends.

Currently, three SSIS versions are supported: SQL Server 2016 Integration Services, SQL Server 2017 Integration Services, and SQL Server 2019 Integration Services.

TencentDB for SQL Server single-node (formerly Basic Edition) and two-node (formerly High Availability/Cluster Edition) instances can use SSIS capabilities through the business intelligence server, while read-only instances cannot.

Currently, the business intelligence server feature is available only in four regions: Guangzhou, Shanghai, Beijing, and Hong Kong (China).

Notes

SSIS capabilities in Tencent Cloud currently can be used for TencentDB for SQL Server and flat files (with .txt, .csv, .xlsx, and .xls extensions).

The business intelligence server and source/target SQL Server instances must be in the same region. SQL Server database instances can communicate with business intelligence servers in different AZs in the same region.

One business intelligence server can be connected to an unlimited number of database instances. In other words, you can connect one business intelligence server to multiple source and target instances to run multiple SSIS project tasks.

The CPU and memory specifications of a business intelligence server are subject to the complexity of SSIS project tasks. The disk space is subject to the size of added flat files.

You can get the **COS File URL** for file upload and deploy the flat file to the business intelligence server only after uploading the flat file to COS. Note that the access permission of the COS object must be set to public read/private write. Only flat files in .txt, .csv, .xlsx, or .xls format are supported. The filename must start with a letter and can contain only digits, letters, underscores, and hyphens.

A domain prefix will be automatically added to Windows authentication accounts of the business intelligence server created in the console, and you don't need to care about the prefix. For example, if you create an account `act1` in the console, the account name displayed in the list will be `xx_x_xx_xxxx/act1`.

Source and target database instances and the business intelligence server involved in an SSIS project need to be interconnected before the project can be deployed. Therefore, all of them need to be added to the same interworking group, and the interworking IP for business intelligence services needs to be enabled for each of them.

After the database instances and business intelligence server are added to an interworking group, each of them has two IPs: private IP and interworking IP for business intelligence services, with different purposes. Therefore, carefully distinguish between them in operation steps.

SSIS projects only support the project deployment mode.

You can use SQL Server Agent to run SSIS program packages.

Do not manually create or restore the `SSISDB` database; otherwise, SSIS may not run properly.

Flowchart

Prerequisites

Prepare a built SSIS project file with the `.ispac` extension.

1. [Purchase a business intelligence server.](#)

TencentDB for SQL Server uses SSIS capabilities, which require project deployment through the SSIS engine in a business intelligence server. If you are using SSIS for the first time, you need to purchase a business intelligence server. If the source and target TencentDB for SQL Server instances in your SSIS project already have a business intelligence server in the same region, skip this step and proceed to **step 2**.

2. [Create a Windows authentication account.](#)

You need to create a Windows authentication account for the business intelligence server for login and SSIS project deployment.

Note:

Accounts created on the business intelligence server all have Windows authentication permissions. Business intelligence servers can use only this type of accounts, and account permissions cannot be modified.

3. [Add a flat file.](#)

Before deploying an SSIS project, you need to check whether the project involves flat files, and if so, you need to add the flat files to the business intelligence server first. If the source and target instances in the project don't involve flat files, skip this step and proceed to **step 4**.

4. [Interconnect the source and target instances and business intelligence server.](#)

Before deploying an SSIS project, you need to interconnect the source and target database instances and business intelligence server involved in the project. The interworking group management feature is used to sustain interconnection between instances. Therefore, in the same account and region, all database instances and the business intelligence server in the SSIS project need to be added to the same interworking group, and the interworking IP for business intelligence services needs to be enabled for each of them before they can access each other.

5. [Deploy an SSIS project.](#)

Before deploying an SSIS project, you need to connect to the business intelligence server as follows:

5.1 Create a Windows authentication account with the same account name and password as that on the business intelligence server in a Windows CVM instance.

5.2 Use the Windows authentication account created in step 5.1 to log in to the Windows CVM instance.

5.3 Use the Windows authentication account to log in to the business intelligence server.

5.4 Deploy an SSIS project.

5.5 Configure the SSIS service, including connection configurations for flat files and the source and target TencentDB for SQL Server instances.

5.6 Run the SSIS service and execute the package.

5.7 Configure an agent job by creating job steps and schedule.

SSIS Operations

Feature Page	Feature	Operation Guide and Directions
Instance List	Purchasing business intelligence server	For detailed directions on how to create/purchase a business intelligence server, see Purchasing Business Intelligence Server .
	Renaming instance	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Renaming Instance .
	Restarting instance	The operation steps are the same as those for a TencentDB for SQL Server instance.

		For detailed directions, see Restarting Instance .
	Terminating instance	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Terminating Instance .
	Recycle bin	After a business intelligence server is terminated or deleted by mistake, the operations that can be performed in the recycle bin are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Recycle Bin .
	Editing tag	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Setting Instance Tag .
Instance management	Setting instance remarks	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Setting Instance Remarks .
	Changing network	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Changing Network (from VPC to VPC) .
	Modifying project	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Setting Instance Project .
	Setting instance maintenance information	The operation steps are the same as those for a TencentDB for SQL Server instance. For detailed directions, see Setting Instance Maintenance Information .
	View monitoring chart	The monitoring metrics of a business intelligence server and a TencentDB for SQL Server instance are not completely the same. Specific monitoring metrics are as displayed on the actual monitoring page. For detailed directions, see Viewing Monitoring Charts .

	Configuring security group	A business intelligence server also needs to be configured with a security group. For detailed directions, see Configuring Security Group .
	Managing account	For detailed directions on how to create and delete accounts, see Creating Windows Authentication Account .
	Managing SSIS	For detailed directions on how to manage SSIS and add flat files, see Adding Flat File .
Interworking Group Management	Interconnecting source/target instances and business intelligence server	For detailed directions on how to enable the interworking IP for business intelligence services and add instances to an interworking group, see Interconnecting Source/Target Instances and Business Intelligence Server .

Purchasing Business Intelligence Server

Last updated : 2024-01-18 17:23:30

SSIS is an important component for SQL Server to implement BI. It provides enterprise-grade data integration and conversion solutions to extract, transform, and load (ETL) data from various sources.

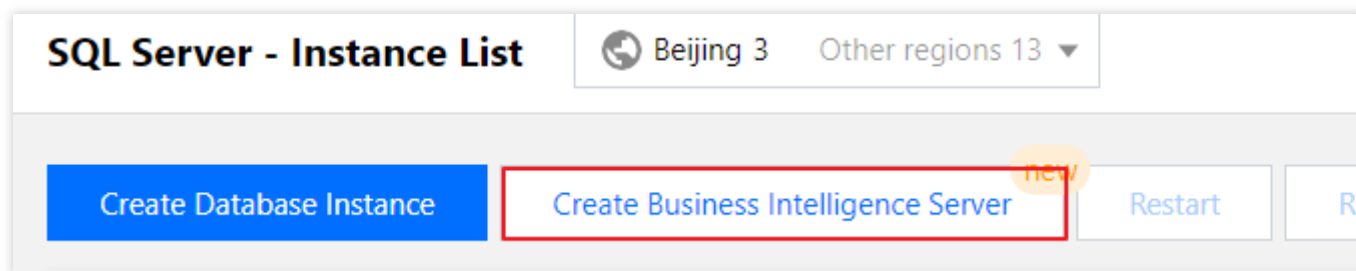
This document describes how to purchase a business intelligence server in the console before you can use the SSIS feature.

Billing

Currently, the business intelligence server feature is in beta test and can be used free of charge. During the beta test, you can purchase only one business intelligence server with the 2-core 4 GB MEM specification in each region and up to three under each root account. Billing will start in pay-as-you-go mode after the beta test ends.

Creating a Business Intelligence Server

1. Log in to the [TencentDB for SQL Server console](#).
2. On the instance list page, click **Create Business Intelligence Server**.



3. On the business intelligence server purchase page, complete the following **configuration item**, carefully read and indicate your content to the Terms of Service, and click **Buy Now**.

Basic Info

Billing Mode

Pay as You Go

Region

South China

East China

North China

Hong Kong/Macao/Taiwan (China Region)

Guangzhou

Tencent Cloud products in different regions can't communicate with each other through private network. Thus, we recommend you select the business intelligence server in the same region as your SQL Server database. For example, the business intelligence server in Beijing only supports SQL Server in Beijing.

AZ

☒ Randomly Assign

Network

4093 subnet IPs in total, with 4076 available

If the existing VPCs and subnets can't meet your requirements, You can [Create VPCs](#) or [Create Subnets](#). After business intelligence server is purchased, the VPC and subnet can be modified in the console.

Intelligence Engine

NEW

Coming Soon

Coming Soon

SSIS

SSAS

SSRS

SSIS can be used to extract, transform, and load (ETL) data from multiple sources to help you solve problems in business intelligence analysis, high-value data mining, and system construction for master data management in an all-round way.

Version

SQL Server 2016 Integration Services

SQL Server 2017 Integration Services

SQL Server 2019 Integration Services

Disk Type

High-Performance Cloud

Specification

2-core 4 GB memory

The specification depends on the computational complexity of the data flow in the SSIS project.

Disk

0 GB

50 GB

100 GB

150 GB

200 GB

20 GB

Disk capacity depends on the size of flat files.

Maintenance Window

☒ Mon
 ☒ Tue
 ☒ Wed
 ☒ Thu
 ☒ Fri
 ☒ Sat
 ☒ Sun

Maintenance Time

Start Time

00:00

Duration

6

Hours

Project List

DEFAULT PROJECT

Security Group

Instruction

To open other ports, you can [Create Security Groups](#)

Tag

+ Add

Quantity

1

Configuration Fees

ICNY

Buy Now

Billing Mode: Pay-as-you-go mode is used by default.

Region: Tencent Cloud products in different regions can't communicate with each other through private network. Thus, we recommend you select the business intelligence server in the same region as your SQL Server database. For example, the business intelligence server in Beijing only supports SQL Server in Beijing.

AZ: You can select **Randomly Assign** or manually select one.

Randomly Assign: **Randomly Assign** is selected by default. The system will select an AZ automatically. AZs in the same region are connected through low-latency private network linkages. We recommend you select **Randomly Assign**.

Manual selection: You can deselect **Randomly Assign** and manually select an AZ supported in the region.

Note:

SQL Server database instances can communicate with business intelligence servers in different AZs in the same region.

Network: You can select a VPC containing the subnet and subnet IP. If the existing VPCs and subnets can't meet your requirements, you can [create VPCs](#) or [create subnets](#).

Intelligence Engine: Currently, you can select only the SSIS intelligence engine. SSAS and SSRS intelligence engines will be available in the future.

Version: Currently, three SSIS versions are supported: SQL Server 2016 Integration Services, SQL Server 2017 Integration Services, and SQL Server 2019 Integration Services.

Disk Type: High-performance cloud disk is used by default.

Specification: Currently, only the 2-core 4GB MEM specification is supported. The specification is subject to the project complexity in the SSIS package.

Disk: Currently, you can select a storage space between 20 and 200 GB. The disk size is subject to the size of flat files to be added.

Maintenance Window and Maintenance Time: To ensure the stability of the business intelligence server, the backend system performs maintenance operations on the instance during the maintenance window from time to time. We recommend you set an acceptable maintenance time for your business instance, usually during off-peak hours, so as to minimize the potential impact on your business. You can manually set the maintenance date and maintenance duration.

Project List: You can assign the instance to different projects for management.

Security Group: It serves as a stateful virtual firewall with filtering feature for configuring network access control for one or more TencentDB instances. It is an important network security isolation tool provided by Tencent Cloud. You can select a security group for the business intelligence server.

Tag: You can set tags for the business intelligence server to categorize and manage resources more easily.

Terms of Service: Read and indicate your consent to the [Terms of Service](#).

4. Return to the instance list after purchase. When the status of the purchased business intelligence server becomes **Running**, the instance is purchased successfully.

Viewing the Purchased Business Intelligence Server

After purchasing the business intelligence server successfully, you can view its information in the instance list in the TencentDB for SQL Server console.

Instance ID/Name	Status	Project	Version	Configuration	Network	AZ
[Redacted]	Running	Default Project	SQL Server 2016 Integration Services	2-core-4GB/20GB Business Intelligence Services - High-Performance Cloud Disk	[Redacted]	Guangzhou Zone 6

Instance ID/Name: It must start with `mssqlbi` , such as `mssqlbi-xxxxxxx` . When you hover over the instance ID, the quick copy button



will be displayed. You can modify the instance name, which can contain up to 60 letters, digits, or underscores.

Status: Status of the business intelligence server, which can be filtered.

Project: Project of the business intelligence server, which can be copied and filtered.

Version: Version of the business intelligence server, which can be filtered.

Configuration: Configuration information of the business intelligence server.

Network: Network of the business intelligence server.

AZ: AZ of the business intelligence server, which can be filtered.

Private Network Address: Private IP and port number of the business intelligence server.

Billing Mode: Billing mode of the business intelligence server.

Tag: Tag keys and values of the business intelligence server.

Operation: Operations that can be performed on the business intelligence server, including **Manage**, **Terminate**, and **Edit Tag**.

You can perform the following operations on the purchased business intelligence server (i.e., SSIS instance) in the instance list:

Feature	Operation
Modify instance name	You can rename a purchased SSIS instance in the console as instructed in Renaming Instance .
Restart instance	You can restart a purchased SSIS instance in the console as instructed in Restarting Instance .
Terminate instance	You can terminate a purchased SSIS instance in the console as instructed in Terminating Instance .
Edit tag	You can modify the tags of a purchased SSIS instance as instructed in Setting Instance Tag .


Creating Windows Authentication Account

Last updated : 2024-01-18 17:23:30


This document describes how to create a Windows authentication account for a business intelligence server, reset the account password, and delete the account in the console.

Creating an account

1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, find the target business intelligence server and click its ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select **Account Management** and click **Create Account**.



4. On the **Create Account** page, set configuration items and click **OK**.

Create Account

Account Name
✓

The account name must contain 1-50 letters, digits, or underscores, and start with a letter.

New Password
!

The password must contain 8–32 characters in at least three of the following four types: uppercase letters, lowercase letters, digits, and symbols (_+-&=!@ \$#%^*()[]).

Confirm Password
✓

Remarks

Enter remarks

Up to 256 characters for remarks

OK

Cancel

Account Name: It can contain 1–50 letters, digits, or underscores and must start with a letter.

New Password: It must contain 8–32 characters in at least three of the following four types: uppercase letters, lowercase letters, digits, and special symbols `_+-&=!@ $#%^*() []`. It cannot contain two or more consecutive characters of the account name.

Confirm Password: Enter the password again.

Remarks: You can enter up to 256 characters.

Note:

Business intelligence servers accounts created in the console all have Windows authentication permissions. Business intelligence servers can use only this type of accounts, and account permissions cannot be modified.

5. After successful creation, you can view the information of the new account on the **Account Management** tab, including account name, status, account creation time, account update time, password update time, and remarks. You can also perform operations on it, including **Reset Password** and **Delete Account**.

Create Account		Batch Management					
<input type="checkbox"/>	Account Name	Status	Account Creation Time	Account Update Time	Password Update Time	Remarks	Operation
<input type="checkbox"/>	sa	Normal	2022-07-04 10:50:59	2022-07-04 10:51:39	2022-07-04 10:51:39	...	Reset Password Delete Account
<input type="checkbox"/>	sa	Normal	2022-07-04 10:51:18	2022-07-04 10:52:00	2022-07-04 10:52:00	...	Reset Password Delete Account

Note:

A domain prefix will be automatically added to business intelligence server accounts created in the console, and you don't need to care about the prefix. For example, if you create an account `act1` in the console, the account name displayed in the list will be `xx_x_xx_xxxx/act1`.

Resetting a password

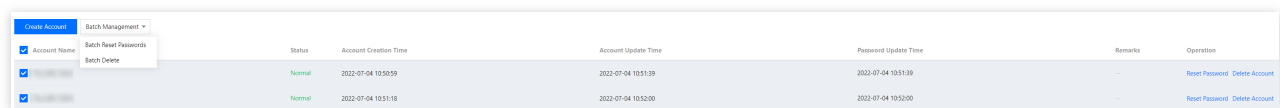
If you forgot the password of a created business intelligence server account, or you need to reset the password, you can reset the password of one or multiple accounts on the **Account Management** tab.

Reset the password of a single account:

Find the target account in the account list, click **Reset Password** in the **Operation** column, enter and confirm the new password, and click **OK**.

Reset the password of multiple accounts:

Select target accounts in the account list, select **Batch Management > Batch Reset Passwords** above the list, enter and confirm the new password, and click **OK**.



The screenshot shows the 'Account Management' tab in the Tencent Cloud console. A dropdown menu is open, showing options: 'Create Account', 'Batch Management', 'Batch Reset Passwords', and 'Batch Delete'. Below the menu is a table with columns: 'Account Name', 'Status', 'Account Creation Time', 'Account Update Time', 'Password Update Time', 'Remarks', and 'Operation'. Two accounts are listed, both with a status of 'Normal'.

Account Name	Status	Account Creation Time	Account Update Time	Password Update Time	Remarks	Operation
xx_x_xx_xxxx/act1	Normal	2022-07-04 10:50:59	2022-07-04 10:51:39	2022-07-04 10:51:39	...	Reset Password Delete Account
xx_x_xx_xxxx/act2	Normal	2022-07-04 10:51:18	2022-07-04 10:52:00	2022-07-04 10:52:00	...	Reset Password Delete Account

Note:

Batch password resetting resets the password of all selected accounts to the same password. To set different passwords for different accounts, you need to reset the password of each account.

Deleting an account

You can delete one or multiple accounts on the **Account Management** tab.

Delete one account

Find the target account in the account list, click **Delete Account** in the **Operation** column, and click **Delete** in the pop-up window.

Batch delete accounts

Select target accounts in the account list, select **Batch Management > Batch Delete** above the list, and click **OK** in the pop-up window.

Note:

To prevent deletion failures, the system will first close all connections to the account before you delete it.

Adding Flat File

Last updated : 2024-01-18 17:23:30

Generally, data is stored in many different data storage systems, and you need to extract data from these data sources, transform it, and load it into one or multiple objects. The SSIS feature of TencentDB for SQL Server supports flat files. If your SSIS project involves a flat file, you need to deploy it to a business intelligence server first before deploying the project.


Note:

If your SSIS project doesn't involve flat files, you can skip this step.

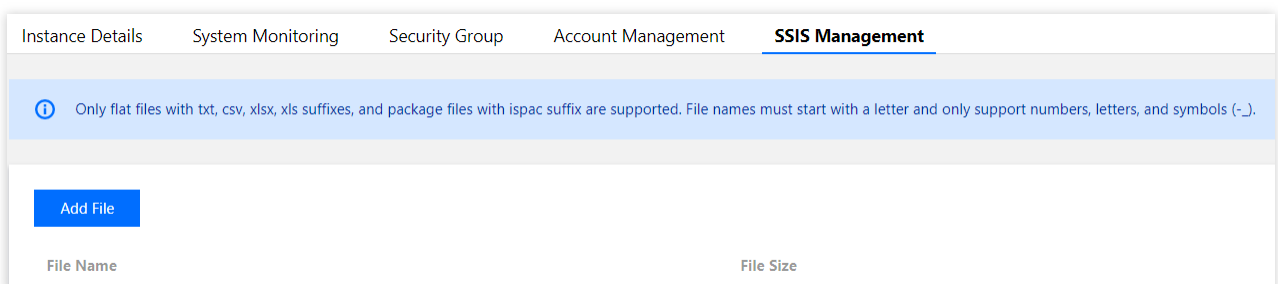
This document describes how to use the SSIS management feature to add and deploy a flat file for easier data use and management.

Directions

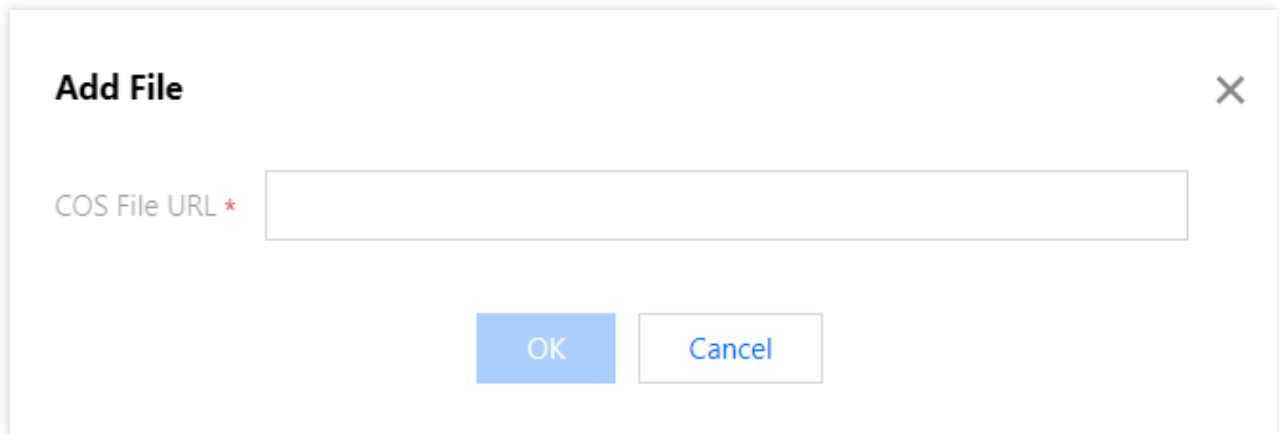
1. Log in to the [TencentDB for SQL Server console](#).
2. In the instance list, find the target business intelligence server and click its ID or **Manage** in the **Operation** column to enter the instance management page.



3. On the instance management page, select **SSIS Management** > **Add File**.



4. In the **Add File** window, paste the **COS File URL** and click **OK**.

A dialog box titled "Add File" with a close button (X) in the top right corner. It contains a label "COS File URL *" followed by a text input field. At the bottom, there are two buttons: "OK" and "Cancel".

Add File ✕

COS File URL *

OK **Cancel**

Note:

You can get the **COS File URL** for file upload and deploy the flat file to the business intelligence server only after uploading the flat file to COS.

You can use COS as instructed in [Downloading Objects](#). The **COS File URL** can be obtained during file download. Note that the access permission of the COS object must be set to **public read/private write**; otherwise, the flat file will fail to be deployed.

Only flat files in .txt, .csv, .xlsx, or .xls format are supported. The filename must start with a letter and can contain only digits, letters, underscores, and hyphens.

5. After adding the flat file successfully, you can query the file information in the list, including the file name, file size, file path, and status. You can also perform operations (**Delete** and **Copy MD5**) on the file.

File Name: Name of the flat file.

File Size: Size of the flat file. A flat file uses the disk space of the business intelligence server.

File Path: Path of the flat file , which is fixed at `D:\\\\SSIS\\\\flat file name` and is required during SSIS project deployment.

Status: Deployment status of the flat file, including **Deployment successful**, **Deploying**, and **Deployment failed**. If the status is **Deployment successful**, the flat file has been deployed to the business intelligence server successfully, and then you can deploy SSIS projects containing the file. If the status is **Deployment failed**, check whether the access permission of the COS object is set to **public read/private write**.

Operation (Delete): You can delete the flat file uploaded to the business intelligence server.

Operation (Copy MD5): You can verify whether the flat file uploaded to the business intelligence server has the same content as the local flat file. To do so, click **Copy MD5** to get the file's MD5 value and compare it with that of the local file.

Interconnecting Source/Target Instances and Business Intelligence Server

Last updated : 2024-01-18 17:23:30

Before using the SSIS feature, you need to interconnect the source instance, target instance, and business intelligence server.

The interworking group management feature is used to sustain interconnection between instances. Therefore, in the same account and region, all database instances and the business intelligence server in an SSIS project need to be added to the same interworking group, and the interworking IP for business intelligence services needs to be enabled for each of them before they can access each other.

This document describes how to add/manage instances in an interworking group.

Adding an Instance to an Interworking Group

1. Log in to the [TencentDB for SQL Server console](#).
2. Select **Interworking Group Management** on the left sidebar, select a region at the top of the page, and click **Add Instance**.
3. In the **Add Instance** window, select the target database instances and business intelligence server, enable interworking IP for business intelligence services for each of them by toggling on the switch, and click **Next**.

Add Instance

1 Enable interworking IP for business intelligence services >

2 Add enabled instance to interworking group

Instance ID/Name	Version	Enable interwork
	SQL Server 2012 Enterprise	<input checked="" type="checkbox"/>
	SQL Server 2012 Enterprise	<input checked="" type="checkbox"/>
	SQL Server 2017 Enterprise	<input type="checkbox"/>

[Cancel](#)[Next](#)

4. Confirm the information of the configured instances and click **OK** to add them to the interworking group.

Add Instance



Enable interworking IP for business intelligence services >



Add enabled instance to interworking group

After you click OK, the following database instances will be added to the interworking group if they have enabled

Instance ID/Name	Version
	SQL Server 2012 Enterprise
	SQL Server 2012 Enterprise

[Previous](#)
[OK](#)

5. Return to the **Interworking Group Management** page. After the status of the instances just added to the interworking group becomes **Added to interworking group**, the instances are successfully added.

After adding instances successfully, you can view the following information in the instance list on the **Interworking Group Management** page.

<input type="checkbox"/> Instance ID/Name	Status	Version	Interworking IP for Business Intelligence Services
<input type="checkbox"/> msq1-	Added to interworking group	SQL Server 2012 Enterprise	172. 024
<input type="checkbox"/> msq2-	Added to interworking group	SQL Server 2012 Enterprise	172. 025

Instance ID/Name: ID/name of the instance. You can rename the instance here.

Status: You can filter instances by interworking status, including **Enabling interworking IP**, **Enabled interworking IP**, **Added to interworking group**, **Reclaiming interworking IP**, and **Reclaimed interworking IP**.

Version: Instance version information, which can be filtered.

Interworking IP for Business Intelligence Services: The interworking IP for business intelligence services is displayed here.

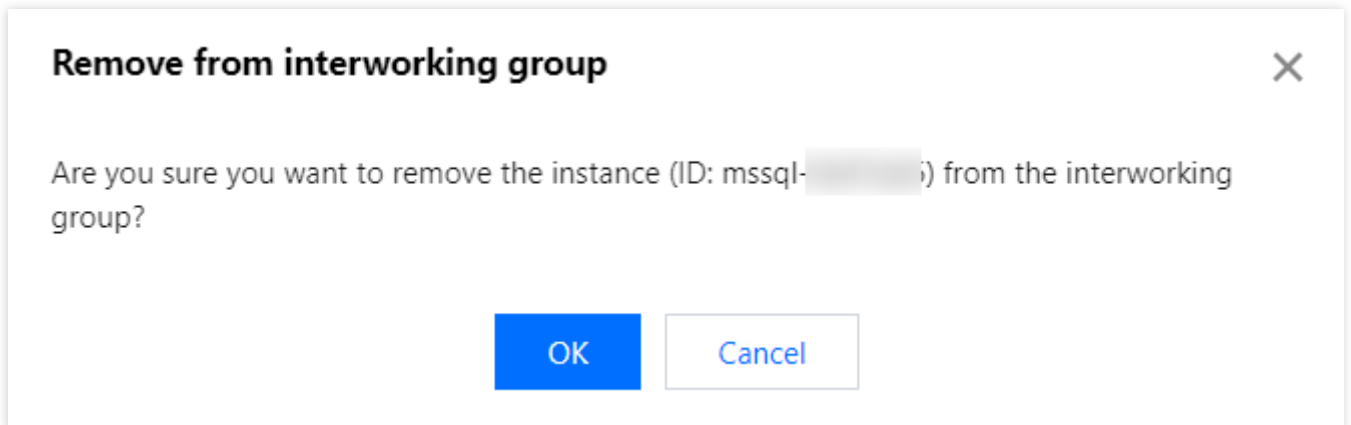
AZ: AZs in the current region, which can be filtered.

Private Network Address: Private network address of the TencentDB for SQL Server instance.

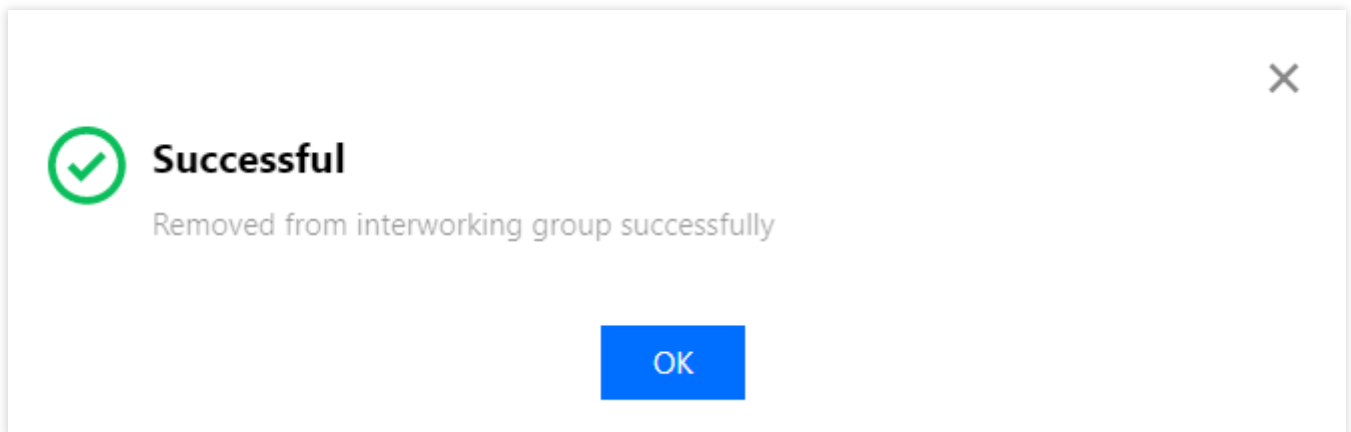
Operation (Remove): You can remove an instance from the interworking group. Once removed, the instance cannot interconnect with other instances in the group.

Removing an Instance from an Interworking Group

1. Log in to the [TencentDB for SQL Server console](#).
2. Select **Interworking Group Management** on the left sidebar and select a region at the top of the page.
3. On the **Interworking Group Management** page, select the target instance in the instance list and click **Remove** in the **Operation** column.
4. In the pop-up window, confirm the selected instance and click **OK**.



5. After removing the instance, click **OK** in the pop-up window, and the removed instance will no longer be displayed in the instance list.



Note:

Once removed, the instance cannot interconnect with other instances in the group.
To make the instance interconnect with other instances in the group, add it back to the group.

Deploying Project

Last updated : 2024-01-18 17:23:30

This document describes how to connect to a business intelligence server and deploy a project through a Windows CVM instance.

Step 1. Create a Windows authentication account with the same account name and password as that on the business intelligence server in a CVM instance

You have created a Windows authentication account on the business intelligence server. Then, you need to create a Windows authentication account with the same account name and password as that on the business intelligence server in a CVM instance. For more information, see [Creating Windows Authentication Account](#).

1. [Log in to the CVM instance](#) and click **Control Panel** in the **Start** menu of Windows Server.
2. In the Control Panel, click **Change account type**.
3. On the **Manage Accounts** page, click **Add a user account**.
4. In the **Add a user** pop-up window, enter a Windows authentication account with the same account name and password as that on the business intelligence server in the console, and click **Next**.
5. Click **Finish**.
6. Return to the **Manage Accounts** page, and you can see the account just created.

Step 2. Use the Windows authentication account created in step 1 to log in to the CVM instance

After creating a Windows authentication account with the same account name and password as that on the business intelligence server, use this account to log in to the CVM instance.

If you are logging in remotely, you need to perform the following operations to add the remote login permission to the account created in the previous step. Otherwise, you can skip this step and directly log in to the CVM instance with this account.

Add the remote login permission to the Windows authentication account created in the CVM instance in the following steps:

1. [Log in to the CVM instance](#) and click **Control Panel** in the **Start** menu of Windows Server.
2. In the search box in the Control Panel, enter **remote**, refresh the page, and click **Allow remote access to your computer**.

3. On the **Remote** tab in the **System Properties** pop-up window, click **Select Users**.
4. In the **Remote Desktop Users** pop-up window, click **Add**.
5. In the **Select Users** pop-up window, enter the Windows authentication account with the same account name and password as that on the business intelligence server created in the CVM instance in **Enter the object names to select** and click **OK**.
6. In the **Remote Desktop Users** window, you can see that the remote login permission is granted to the account. Click **OK**.

Step 3. Use the Windows authentication account to log in to the business intelligence server

1. Log in to the [TencentDB for SQL Server console](#) and view the private network address of the business intelligence server in the instance list.
2. Download and install [SSMS](#) in the CVM instance. For more information on SSMS, see [What is SQL Server Management Studio \(SSMS\)?](#).
3. Start SSMS in the Windows CVM instance.

In the **Connect to Server** window, enter the relevant information to connect to the business intelligence server, click **Connect**, and wait a few minutes while SSMS connects to your business intelligence server through Windows authentication.

Server type: Select **Database Engine**.

Server name: Enter the private IP and port number of the business intelligence server and separate them by a comma.

For example, if the private IP of the instance is `10.10.10.10` and the port number is `1433`, then enter

`10.10.10.10,1433` here.

Authentication: Select **Windows Authentication**.

User name and password: You don't need to enter anything, as the information of the created Windows authentication account will be entered automatically. A prefix will be automatically added to the account name here, so that it can be the same as that displayed in the list on the **Account Management** tab of the instance.

4. After logging in to SSMS, you can see the private IP of the business intelligence server in the **Object Explorer**, indicating that you have successfully logged in to the instance with the Windows authentication account.

Step 4. Deploy an SSIS project

1. View the SSISDB.

In the **Object Explorer** of SSMS, expand the **Integration Services Catalogs** directory to view the SSISDB.

2. Create a folder.

Right-click **SSISDB** and click **Create Folder**. In the pop-up window, enter the folder name and description and click

OK.

3. Expand the folder created in the above step, and you can see the **Projects** and **Environments** directories.

4. Deploy the project. Right-click **Projects** and click **Deploy Project**

5. On the **Introduction** page in the **Integration Services Deployment Wizard**, click **Next**.

6. Select an SSIS project deployment source.

On the **Select Source** page in the **Integration Services Deployment Wizard**, select **Project Deployment > Project deployment file**, click **Browse**, select the path of the prepared SSIS project file to be deployed, i.e., the .ispac file, and click **Next**.

Note:

SSIS projects only support the project deployment mode.

7. During deployment, the **SQL Server Integration Services** warning window may pop up. If it does, click **OK** to ignore it.

8. Select the deployment target.

On the **Select Deployment Target** page in the **Integration Services Deployment Wizard**, click **SSIS in SQL Server** and **Next**.

9. Select the destination.

On the **Select Destination** page in the **Integration Services Deployment Wizard**, **Server name** is the business intelligence server's private IP and port, and **Authentication** is **Windows Authentication** by default. Leave the settings as-is and click **Connect**.

10. After clicking **Connect**, you can see that the **Path** field is activated, and the .ispac SSIS project file is displayed after the path of the folder created previously. At this point, directly click **Next**.

11. Check the SSIS project deployment options. On the **Review** page in the **Integration Services Deployment Wizard**, check whether the source and destination information of the SSIS project to be deployed is correct. After confirming that everything is correct, click **Deploy** to deploy the project in the created folder.

12. On the **Results** page in the **Integration Services Deployment Wizard**, if all results are **Passed**, the SSIS service is deployed successfully.

Step 5. Configure the SSIS service

1. After the SSIS project is deployed, configure relevant services.

After deployment, you can view the .ispac SSIS project file in the created **Projects** directory.

2. Right-click the deployed SSIS project file and click **Configure**.

3. In the **Configure** pop-up window, switch to the **Connection Managers** tab.

4. Configure the flat file connection.

5. If the SSIS project file involves a flat file, click Flat File Connection Manager on the Connection Managers tab, find **ConnectionString** in Properties on the right, and click ... after it.

Note:

This step configures the connection to the flat file. If the SSIS project file involves a flat file, perform this step; otherwise, skip it.

6. In the **Set Parameter Value** pop-up window, click **Edit value** in the **Value** section.
7. Log in to the [TencentDB for SQL Server console](#), click the ID of the business intelligence server in the instance list to enter its SSIS management page, and view and copy the **File Path** of the flat file.
8. In the **Set Parameter Value** pop-up window, paste the copied flat file path in the input box after **Edit value** and click **OK**.
9. Return to **Connection Managers**. You can see that the **ConnectionString** property of the **Flat File Connection Manager** has been changed to the path of the flat file uploaded on the business intelligence server's SSIS management page.
10. Configure the connection addresses of the source and target TencentDB for SQL Server instances in the SSIS project file.

On the **Connection Managers** tab, you can see the source and target instance connectors named **TencentDB for SQL Server instance's private IP,port.database name.account name**.

11. Set the connection configuration of the source database first. Click the name of the source database connector, find **ServerName** in **Properties** on the right, and click ... after it.
12. In the **Set Parameter Value** pop-up window, click **Edit value** in the **Value** section.
13. Log in to the TencentDB for SQL Server console. On the **Interworking Group Management** page, find the source TencentDB for SQL Server database instance and copy its interworking IP and port number for business intelligence services.
14. In the **Set Parameter Value** pop-up window, paste the copied interworking IP and port number for business intelligence services of the source TencentDB for SQL Server instance in the input box after **Edit value**, separate the IP and port number by a comma, and click **OK**.

For example, if the interworking IP for business intelligence services is `10.10.10.10` and the port number is `1024`, then enter `10.10.10.10,1024` here.

15. Return to **Connection Managers**. You can see that the connection address in the **ServerName** property of the **source TencentDB for SQL Server instance** has been changed to the interworking IP for business intelligence services of the source instance.
16. On the **Connection Managers** tab, click the name of the source database connector, find the **Password** property of the **source TencentDB for SQL Server instance** in **Properties** on the right, and click ... after it.
17. In the **Set Parameter Value** pop-up window, select **Edit value** in the **Value** section, enter the password of the source database account in the input box after **Edit value**, and click **OK**.
18. Return to **Connection Managers**, configure the connection information of the target database, including the **ConnectionString** and **Password** properties, and click **OK**.

Step 6. Run the SSIS service

1. After configuring relevant services of the SSIS project file, in the created **Projects** directory, you can view the .dtsx SSIS package file in the **Packages** directory of the SSIS project file.
2. Right-click the .dtsx SSIS package file and click **Execute**.
3. In the **Execute Package** pop-up window, check the configuration information on the **Connection Managers** tab and click **OK**.
4. During execution, the **Microsoft SQL Server Management Studio** prompt window may pop up. If it does, click **Yes**.
5. View the report. If all results are **Succeeded**, the execution succeeds.

Step 7. Configure an agent job

1. In the **Object Explorer** of SSMS, expand **SQL Server Agent**, right-click **Jobs**, and click **New Job**.
2. In the **New Job** pop-up window, click ... after **Owner**. In the **Select Login** pop-up window, enter the **business intelligence server's Windows authentication account** in the input box below **Enter the object names to select**, and click **Check Names**.
3. As you cannot directly use the business intelligence server's default Windows authentication account, the **Multiple Objects Found** window will pop up. Select matched object and click **OK**.
4. In the **Select Login** pop-up window, you can see that the real domain prefix is added before the account name in the input box below **Enter the object names to select**. Click **OK**.

The displayed account name is the same as that displayed in the list on the **Account Management** tab of the business intelligence server.

5. Create a job.

In the **New Job** pop-up window, click **Steps** below **Select a page**. On the **Steps** page, click **Create**, and the **New Job Step** window pops up.

6. Create a job step.

In the **New Job Step** pop-up window, enter **Step name**, **Type**, **Run as**, **Server**, and **Package** as follows:

Step name: Enter a custom job step name.

Type: Select **SQL Server Integration Services Package** in the drop-down list.

Run as: Select the corresponding role of the account in the drop-down list.

Server: Enter the private IP and port number of the business intelligence server and separate them by a comma.

Package: Click ... after the package. In the **Select an SSIS Package** pop-up window, select the SSIS package to be configured with the agent job and click **OK**.

Note:

Log in to the [TencentDB for SQL Server console](#). In the instance list, find the business intelligence server, and you can view its private network address in the **Private Network Address** column.

7. After selecting the SSIS package to be configured with the agent job, return to the **New Job Step** pop-up window.

8. In the **New Job Step** pop-up window, set **Server** to the interworking IP and port number for business intelligence services, separate the IP and port number by comma, and click **OK**.

Note:

Log in to the [TencentDB for SQL Server console](#). On the **Interworking Group Management** tab, find the business intelligence server, and you can view its interworking IP for business intelligence services.

9. At this point, if you click **OK**, the program will stop responding temporarily, and the **SSIS Execution Properties** window will pop up to report a connection error. You can click **OK** to ignore this error.

10. Set the scheduling cycle.

In the **New Job** pop-up window, click **Schedule** below **Select a page**. On the **Schedule** page, click **Create**, and the **New Job Schedule** window will pop up.

11. In the **New Job Schedule** window, set the job scheduling cycle of the SSIS package, including **Name**, **Schedule type**, **Frequency**, **Daily frequency**, and **Duration**, based on your business requirements, and click **OK**.

12. Return to the **New Job** pop-up window. At this point, you have configured a job step and schedule. Click **OK**.

13. In the **Object Explorer** of SSMS, expand **SQL Server Agent**, find **Jobs**, and you can see the job just created.

14. Right-click the name of the job just created and click **View History**. In the **Log File Viewer** pop-up window, you can view the job execution history.

15. View job execution logs in the **Log File Viewer**, and you can see that the job is running normally.

16. In SSMS, log in to the target instance. You can see that the data is being extracted, transformed, and loaded into the target instance according to the settings in the SSIS project.

Data Migration (New)

Data Migration Solution Overview

Last updated : 2024-07-31 09:49:41

To meet different needs for on-cloud business or cloud migration business, TencentDB for SQL Server provides corresponding data migration solutions for the following migration scenarios, allowing users to smoothly migrate to TencentDB for SQL Server without affecting business operations.

Overview of Data Migration Solutions

Migration Scenario	Related Operations
Migrating self-built SQL Server databases (IDC self-built and CVM self-built) to TencentDB for SQL Server	Using DTS with either "Full" or "Full + Incremental" selected for hot database migration. For operations, please refer to Migration Operation Guide .
	Using backup and restore for cold data migration. For specific operations, please refer to Cold Backup Migration .
Migrating SQL Server databases from third-party manufacturers (Alibaba Cloud, Huawei Cloud, and AWS) to TencentDB for SQL Server	Using backup and restore for cold data migration. For specific operations, please refer to Cold Backup Migration .
Database migration between TencentDB for SQL Server instances	Using DTS with either "Full" or "Full + Incremental" selected for hot database migration. For operations, please refer to Migration Operation Guide .
	Using DTS for cross-account data migration (suitable for instance migration under different accounts). For specific operations, please refer to Cloud Database Cross-account Migration Guide .
	Using backup and restore for cold data migration. For specific operations, please refer to Cold Backup Migration .

Cold Backup Migration

Last updated : 2024-07-31 10:08:13

This document describes how to migrate data from SQL Server databases in other cloud vendors and self-built SQL Server databases to TencentDB for SQL Server in the TencentDB for SQL Server console.

Overview

Cold backup migration restores data from BAK, TAR, or ZIP files. It is applicable to data migration from SQL Server databases in other cloud vendors and self-built SQL Server databases to TencentDB for SQL Server, and backup can be performed when the system is shut down.

TencentDB for SQL Server supports migrating data to TencentDB for SQL Server via [Cloud Object Storage \(COS\)](#) files or locally uploaded files.

Note:

To restore backups with ZIP and TAR files, you need to ensure that the decompressed files are in the same path folder as the ZIP and TAR files before decompression. If they are in the path folder of the next layer, they cannot be parsed.

Note

Before migration, make sure that the version of the target SQL Server instance is not earlier than that of the source instance.

The name of the migrated database cannot be the same as that of the TencentDB for SQL Server instance.

If the backup file is of considerable size, it is recommended to employ the method of "Download File from COS", which entails initially uploading the file to COS and then using a COS link for restoration. For instructions on uploading backup files to COS and obtaining a COS link, please refer to [Uploading Backup to COS](#).


Full backups + logs restoration: The target instance and source instance don't have to be on the same version, as long as the instance version is high.

Full backups + differential backups restoration: The target instance and the source instance must be on the same version.

You need to pay attention to the case sensitivity of the .bak, .tar, and .zip file extensions. Only lowercase letters are supported, and uppercase letters are not supported.

Directions

1. Log in to the [TencentDB for SQL Server console](#). In the instance list, click an instance ID to access the instance management page.
2. On the instance management page, select the **Backup and Restoration** tab and click **Create**.



3. On the backup restoration creation tab, complete **Backup Restoration Settings**, and click **Create Task**.

1 Backup Restoration Settings

2 Upload Backup File

Task Settings

Task Name *

It can contain up to 60 characters

The task name can contain up to 60 letters, digits, or underscores.

Backup Upload Method *

Upload File

Download File from COS

Sufficient free-of-charge resources

Restoration Mode *

☒ Full Backup File
 ☐ Full backups + Logs
 ☐ Full backups + Differential backups

Target Database Info

Version *

SQL Server 2008 R2 Enterprise

Region *

Beijing/Beijing Zone 4

Instance ID/Name *

mssql-o68z1td8/53328b-2-159622b787ee

Private Network Address *

172.1433

Cancel

Create Task

Parameter	Description
Task Name	The task name can contain up to 60 letters, digits, or underscores.
Backup Upload Method	Local upload and download from COS.
Restoration Mode	Full backups, full backups + logs, full backups + differential backups.

Note:

To cancel the creation of backup restoration task, click **Cancel**, and the task record will not be retained in the backup restoration task list.

4. On the backup restoration creation tab, upload the backup file, and click **Save**.

Backup upload method falls into the following two types based on step 3.

Note :

If the backup file is of considerable size, it is recommended to employ the method of "Download File from COS", which entails initially uploading the file to COS and then using a COS link for restoration. For instructions on uploading backup files to COS and obtaining a COS link, please refer to [Uploading Backup to COS](#).

Upload backup file directly

Download file from COS

Backup Restoration Settings

2 Upload Backup File

Download File from COS

Restoration Mode *

Full Backup File [Get Backup Command](#)

Upload Backup *

Upload

File	Size	Operation
Upload backup file		

Rename Database ⓘ

☒ Enable Database Renaming

Database Name in Backup Files

Restored Database Name

Delete

Add

Database name is required

Download Settings

☐ Auto-restoration upon upload completion
 ☒ Start restoration manually The uploaded file will only be retained for 24 hours

Previous

Save

Parameter	Description
Restoration Mode	<p>When you select full backups + logs or full backups + differential backups, there will be requirements for file names. You can click Get Backup Command to generate the corresponding backup command to generate the backup files that conforms to the file naming conventions.</p> <p>Full backups restoration: You cannot use the full backup name of "full backup + differential backup" or "full backup + log files".</p> <p>Full backups + differential backups restoration:</p>

	<p>Requirements for full backup file name: dbname_localtime_1full1_1noreconvery1.bak</p> <p>Requirements for incremental backup file name: dbname_localtime_1diff1_1noreconvery1.bak</p> <p>Requirements for the name of the last incremental backup file: dbname_localtime_1diff1_1reconvery1.bak</p> <p>Full backups + logs restoration</p> <p>Requirements for full backup file name: dbname_localtime_2full2_2noreconvery2.bak</p> <p>Requirements for log name: dbname_localtime_2log2_2noreconvery2.bak</p> <p>Requirements for the name of the last log: dbname_localtime_2log2_2reconvery2.bak</p>
Upload Backup	Click Upload to import the backup file locally
Rename Database	<p>It is optional. Once enabled, the original database name in the backup file will be reset and then designated as the new database name after it is restored to the cloud database, and you need to enter these two names.</p> <p>Note:</p> <p>Up to 5 databases can be renamed.</p> <p>If a database in the original database was not renamed, its name will remain unchanged after the backup restoration task is completed.</p>
Download Settings	<p>Support Auto-restoration upon upload completion and *Start restoration manually.</p> <p>Auto-restoration upon upload completion: The restoration task will be started immediately after you click Save.</p> <p>Start restoration manually: Click Save, manually start the task on Backup and Restoration** page, and the uploaded backup files are only saved for 24 hours.</p>

Note:

On the backup file upload page, you can click **Previous** to review and edit the backup restoration settings.

If no more operation is needed, return to backup restoration settings page and click **Next** to upload the backup file. To cancel the task, click **Cancel**.

1 Backup Restoration Settings

>

2 Upload Backup File

Task Settings

Task Name *

test

The name consists of Chinese and English characters, numbers, and special characters, with a maximum length of 60 characters.

Backup Upload Method *

Upload File

Download File from COS

Sufficient free-of-charge resources [🔗](#)

Restoration Mode *

☒ Full Backup File

☐ Full backups + Logs

☐ Full backups + Differential backups

Target Database Info

Version *

SQL Server 2008 R2 Enterprise

Region *

Beijing/Beijing Zone 4

Instance ID/Name *

mssql-o68z1td8/53328b6f-xxxx-xxxx-xxxx-xxxx

Private Network Address *

172.x.x.x :1433

Cancel

Create Task

After returning to the backup restoration settings page, you can click **Edit** to edit the task name, backup upload method, and restoration mode. After you click **Edit**, the content of the backup file upload page in step 2 will be cleared, but you can click **Create Task** to enter the next step for resetting.

Backup Restoration Settings

2 Upload Backup File

Download File from COS

Restoration Mode * Full Backup File [Get Backup Command](#)

Upload Backup * Upload

File	Size	Operation
Upload backup file		

Rename Database ⓘ ☒ Enable Database Renaming

Database Name in Backup Files	Restored Database Name	
<input type="text"/>	<input type="text"/>	Delete
<div>Add</div>		

Database name is required

Download Settings

☐ Auto-restoration upon upload completion
☒ Start restoration manually The uploaded file will only be retained for 24 hours

Previous

Save

Parameter	Description
Restoration Mode	<p>When you select full backups + logs or full backups + differential backups, there will be requirements for file names. You can click Get Backup Command to generate the corresponding backup command to generate the backup files that conforms to the file naming conventions.</p> <p>Full backups restoration: You cannot use the full backup name of "full backup + differential backup" or "full backup + log files".</p> <p>Full backups + differential backups restoration:</p> <p>Requirements for full backup file name: dbname_localtime_1full1_1noreconvery1.bak</p> <p>Requirements for incremental backup file name: dbname_localtime_1diff1_1noreconvery1.bak</p> <p>Requirements for the name of the last incremental backup file: dbname_localtime_1diff1_1reconvery1.bak</p> <p>Full backups + logs restoration</p> <p>Requirements for full backup file name: dbname_localtime_2full2_2noreconvery2.bak</p> <p>Requirements for log name: dbname_localtime_2log2_2noreconvery2.bak</p> <p>Requirements for the name of the last log: dbname_localtime_2log2_2reconvery2.bak</p>
COS Source File Link	<p>Paste the source COS backup file link</p> <p>Note:</p>

	<p>When downloading files from COS, you need to pay attention to the permissions of the original COS file link. The following two permissions are supported:</p> <p>Public read permission: Copy the download address.</p> <p>Private read permission: Download via a pre-signed URL in COS. For more information, see Download via Pre-Signed URL.</p> <p>You need to pay attention to the case sensitivity of the .bak, .tar, and .zip file extensions. Only lowercase letters are supported, and uppercase letters are not supported.</p>
Rename Database	<p>It is optional. Once enabled, the original database name in the backup file will be reset and then designated as the new database name after it is restored to the cloud database, and you need to enter these two names.</p> <p>Note:</p> <p>Up to 5 databases can be renamed.</p> <p>If a database in the original database was not renamed, its name will remain unchanged after the backup restoration task is completed.</p>

Note:

On the backup file upload page, you can click **Previous** to review and edit the backup restoration settings.

If no more operation is needed, return to backup restoration settings page and click **Next** to upload the backup file. To cancel the task, click **Cancel**.

1 Backup Restoration Settings

2 Upload Backup File

Task Settings

Task Name *

test

The name consists of Chinese and English characters, numbers, and special characters, with a maximum length of 60 characters.

Backup Upload Method *

Upload File

Download File from COS

Sufficient free-of-charge resources [🔗](#)

Restoration Mode *

☒ Full Backup File

☐ Full backups + Logs

☐ Full backups + Differential backups

Target Database Info

Version *

SQL Server 2008 R2 Enterprise

Region *

Beijing/Beijing Zone 4

Instance ID/Name *

mssql-o68z1td8/53328b6f

Private Network Address *

172.17.0.1:1433

Cancel

Create Task

After returning to the backup restoration settings page, you can click **Edit** to edit the task name, backup upload method, and restoration mode. After you click **Edit**, the content of the backup file upload page in step 2 will be cleared, but you can click **Create Task** to enter the next step for resetting.

5. View backup restoration task in the task list

Full backups restoration: The restoration task will automatically end after it is completed.

Full backups + differential backups and full backups + logs: After full backup restoration is completed, you can also execute a subtask of uploading differential backups/logs. When the last file is uploaded successfully, the entire backup restoration task will automatically end.

Practical Tutorial

TencentDB for SQL Server supports cross-account backup restoration, that is, you can get the backup file download URL of an instance under account A, and then restore data under account B. The following describes how to do this.

Note:

After getting the backup file download URL of an instance under account A, perform backup restoration under account B. The two accounts need to be in the same region.

This backup restoration mode currently supports backup files in .bak, .tar, and .zip formats.

You need to pay attention to the case sensitivity of the .bak, .tar, and .zip file extensions. Only lowercase letters are supported, and uppercase letters are not supported.

To use intra-region cross-account backup restoration, the form of the backup file obtained in the console must be unarchived file.

Currently, the backup file download URL needs to be [decoded](#) as detailed [below].

Directions

1. Log in to the [TencentDB for SQL Server console](#) with account A, and click an instance ID in the instance list to enter the instance management page.
2. On the instance management page, select the **Backup Management** tab and click **View Details** in the **Operation** column of the target unarchived file in the data backup list.

Manual Backup
Auto-Backup Settings
Backup Task Settings
Cross-Region Backup Settings NEW

Data Backup List
Log Backup List

Last hour
Last 24 hours
Today
Yesterday
Last 7 days
Last 30 days
Separate k

2023-02-13 ~ 2023-02-13

Backup Start Time ⓘ	Backup End Time ⓘ	Backup Name	Backup Policy	Backup Mode	Upload Backu...	Backup Size	Database
2023-02-13 03:23:31	2023-02-13 03:23:36	mssql- [REDACTED]	Instance backup	Scheduled back...	Unarchived files	169KB	test [REDACTED]

3. Click **Download** in the pop-up window and copy the download URL.

Download backup file
×

ⓘ Notes:

- 1. To ensure data security, the address is valid for 15 minutes. Upon its expiration, refresh the page to get a new one.
- 2. URL must be enclosed with quotation marks when wget is used to download.

Download Address

https://sqlserver-bucket-bj-1258415541.cos.ap-beijir
Copy

Download
Cancel

4. Click [here](#) to decode the copied URL on the page redirected to.

Note:

Only the first part of the download URL needs to be decoded.

Example: Suppose a download URL is:

```
https://sqlserver-bucket-bj-1258415541.cos.ap-beijing.myqcloud.com/1312368346%2fsqlserver%2fmssql-8e5hjiaiq%2fbackup%2fautoed_instance_58013012_AdventureWorksDW2012_2022_12_29023915.bak?*****
```

You only need to select the URL part that ends with `.bak?` for decoding. Then, combine the output string with the other part of the URL to get the final URL for cross-account backup restoration.

For the above sample URL, take the following part for decoding:

```
https://sqlserver-bucket-bj-1258415541.cos.ap-beijing.myqcloud.com/1312368346%2fsqlserver%2fmssql-8e5hjiaiq%2fbackup%2fautoed_instance_58013012_AdventureWorksDW2012_2022_12_29023915.b
```

ak?

It is decoded to:


```
https://sqlserver-bucket-bj-1258415541.cos.ap-  
beijing.myqcloud.com/1312368346/sqlserver/mssql-  
8e5hjaiq/backup/autoed_instance_58013012_AdventureWorksDW2012_2022_12_29023915.bak?
```

Then, the final URL for cross-account backup restoration is as follows:

```
https://sqlserver-bucket-bj-1258415541.cos.ap-  
beijing.myqcloud.com/1312368346/sqlserver/mssql-  
8e5hjaiq/backup/autoed_instance_58013012_AdventureWorksDW2012_2022_12_29023915.bak?  
*****
```

5. Copy the URL for cross-account backup restoration, log in to the [TencentDB for SQL Server console](#) with account B, and click an instance ID in the instance list to enter the instance management page.

6. On the instance management page, select the **Backup and Restoration** and click **Create**.



7. In the pop-up window, configure the following items and click **Create Task**.

1 Backup Restoration Settings

>

2 Upload Backup File

Task Settings

Task Name *

The task name can contain up to 60 letters, digits, or underscores.

Backup Upload Method *

Upload File

Download File from COS

Restoration Mode *

☒ Full Backup File

☐ Full backups + Logs

☐ Full backups + Differential backups

Target Database Info

Version *

Region * Beijing/Beijing Zone 4

Instance ID/Name * mssql-o68z1td8/53328-59622b787ee

Private Network Address * 172-1433

Cancel

Create Task

Parameter	Description
Task Name	Enter the task name, which can contain up to 60 letters, digits, or underscores.
Backup Upload Method	Select Download File from COS .
Restoration Mode	Select Full Backup File .

8. In the **Upload Backup File** window, paste the URL and click **Save**.

✓ Backup Restoration Settings

>

2 Upload Backup File

Download File from COS

Restoration Mode * Full Backup File [Get Backup Command](#)

COS Source File Link *

Rename Database ⓘ ☐ Enable Database Renaming

Previous

Save

9. Go to the **Backup and Restoration** tab, find the backup task you just created, and click **Start** in the **Operation** column.

10. In the backup and restoration task list, check the migration task status. When it becomes **Migration succeeded**, the cross-account backup restoration is completed.

Using DTS to Migrate Data

Instructions

Last updated : 2024-07-30 17:48:31

Category	Description
Source/Target Type	<p>1. Source Type: Self-built databases (IDC self-built and CVM self-built) SQL Server 2008R2, 2012, 2014, 2016, 2017, 2019, and 2022. TencentDB (intra-account and cross-account) for SQL Server 2008R2, 2012, 2014, 2016, 2017, 2019, and 2022.</p> <p>2. Target Type: TencentDB (intra-account and cross-account) for SQL Server 2008R2, 2012, 2014, 2016, 2017, 2019, and 2022.</p> <p>3. Data Transfer Service (DTS) migration is not supported in network shared storage disk environments.</p>
Cross-region Migration	Currently, cross-region migration is supported in the Chinese Mainland and Hong Kong (China) regions. Other regions do not support cross-region migration.
Migration Object	<p>1. Only database-level migration is supported, that is, all objects in the database must be migrated together. Single-table migration is not supported.</p> <p>2. The migration of basic database and table objects is supported, but migration of instance-level jobs, triggers, DB links (link servers), or user permissions is not supported. These objects need to be rebuilt after the migration is complete..</p>
Source Database Impact	<p>1. During full data migration, 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 perform the migration during off-peak hours.</p> <p>2. The full data migration process is implemented through locked migration, and write operations will be briefly (in seconds) blocked during table locking.</p> <p>3. Due to changes in the physical environment of data files after migration, database indexes will become invalid. You need to rebuild the indexes after migration; otherwise, database performance may significantly decline.</p> <p>4. By default, a lock-free synchronization method is used. In the full data export stage, a global lock (FTWRL) will not be applied on the source database; only table locks are applied on tables without a primary key.</p>
Supported SQL Statements	<p>1. DDL TABLE: CREATE TABLE, ALTER TABLE, DROP TABLE, TRUNCATE TABLE, and RENAME TABLE VIEW: CREATE VIEW, ALTER VIEW, and DROP VIEW INDEX: CREATE INDEX, and DROP INDEX</p>

	<p>DATABASE: CREATE DATABASE, ALTER DATABASE, and DROP DATABASE</p> <p>2. DML</p> <p>INSERT, UPDATE, DELETE, and REPLACE</p>
Limits on Operations	<ol style="list-style-type: none">1. Only one migration task can be initiated at any time for the same source instance.2. Do not modify or delete user information (including username, password, and permissions) in the source and target databases and port numbers during migration; otherwise, the migration task will fail.3. Do not perform transaction log backup during incremental sync; otherwise, the transaction log will be truncated and become discontinuous.4. If you only perform full data migration, do not write new data into the source database during migration; otherwise, the data in the source and target databases will be inconsistent. In scenarios with data writes, to ensure data consistency in real time, we recommend that you select full + incremental data migration.5. For full + incremental data migration, when the task status is completing after clicking Complete, do not write new data into the source database. We recommend stopping data writes for two minutes after clicking Complete; otherwise, the data in the source and target databases may be inconsistent.

Migration Operation Guide

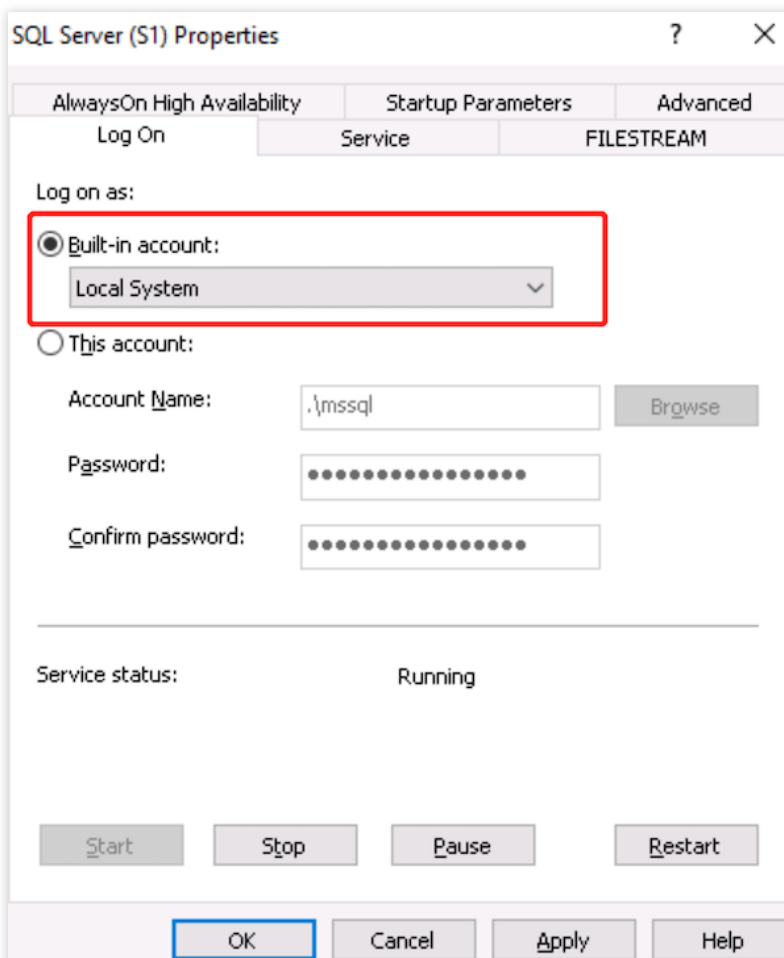
Last updated : 2024-07-30 18:00:41

Operation scenarios

This document provides operation guidance for migrating data from SQL Server to TencentDB for SQL Server using the DTS data migration feature.

Preparations

1. Please carefully read the [Usage Instructions](#) to understand the feature constraints and precautions.
2. In advance, ensure the access path between DTS and the database is established according to the access type you need. For details, refer to [Network Preparation Work](#).
IDC self-built database: You can choose "Public Network/Direct Connect/VPN Access/Cloud Connect Network (CCN)" as the access type.
Self-built databases on cloud virtual machine (CVM): Choose "Self-Built on CVM" as the access type.
TencentDB instances: Select "TencentDB" as the access type.
3. The service where the source database is located must have the file-sharing port 445 enabled.
4. The source database must be set to "Full Recovery Mode", and it is recommended to make a full backup before migration.
5. The local disk space of the source database must be large enough to accommodate the size of the database to be migrated.
6. When the source instance is a non-TencentDB for SQL Server instance (self-built instance on Public Network/CVM) or a TencentDB for SQL Server Basic Edition instance, the target end must use an account with sysadmin permissions for migration and be able to run the xp_cmdshell stored procedure. When the source instance is a TencentDB for SQL Server High Availability Edition or Cluster Edition instance, there are no permission restrictions on the target end account.
7. The SQL service startup account on the migration source needs to be changed to the built-in account Local System. There are no restrictions on the account of the source database to be migrated, but it must have sysadmin permissions.



As shown in the figure, start the SQL service on the migration source. In the startup configuration, select **Log on as**, select **Built-in account**, and change it to Local System startup.

Note:

After modifying the account, you need to restart the SQL server service.

Environment Requirements

Note:

The following environment requirements will be automatically checked by the system before starting the migration task. The system will report an error for requirements not met. If users can identify the issue, they can refer to [Verification Item Check Requirements](#) to make necessary modifications by themselves. If not, wait for the system check to complete and then follow the error prompts to make the necessary modifications.

Type	Environmental Requirements
Source Database Requirements	The service where the source instance is located needs to have the file-sharing port 445 enabled. The source and target databases' networks must be interconnected.

	The server where the source database is located must have sufficient outbound bandwidth; otherwise, the migration rate will be affected.
Target Database Requirements	<p>Only migration from Basic Edition to High Availability Edition (including Dual-Server High Availability Edition and Cluster Edition) is supported, and the version number of the target instance must be later than that of the source instance.</p> <p>The target database cannot have databases with the same name as those in the source database.</p> <p>The disk space of the target database must be larger than the size of the source database, specifically 1.5 times the size of the source database.</p> <p>The target database cannot have access requests or active businesses; otherwise, the migration will fail.</p>

Migration Operation

1. Log in to [DTS console](#), select **Data Migration** on the left navigation bar, and click **Create Migration Task** to enter the Create Migration Task page.
2. On the Create Migration Task page, select the source instance type and region, the target instance type, region, specification, etc., and then click **Buy Now**.

Configuration Item	Description
Source instance type	Please select according to your source database type. Once purchased, it cannot be modified. For this scenario, select "SQL Server".
Source instance region	Select the source database region. If the source database is a self-built one, select a region nearest to it.
Target instance type	Please select according to your target database type. Once purchased, it cannot be modified. For this scenario, select "SQL Server".
Target instance region	Select the target database region.
Specification	Currently, only the fixed specification Medium is supported.

3. 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 is passed, click **Create**.

Note:

If the connectivity test fails, troubleshoot and fix the issues as prompted in the [Repair Guidance](#), then try again.

Settings Type	Configuration Item	Description
Task configuration	Task name	Set a business-significant name for easy task identification.
	Running node	Execute immediately: Starts the task immediately after task validation is passed. Scheduled execution: A task execution time must be configured, and the task will start when the time is reached.
	Tag	Tags are used to manage resources by category from different dimensions. If the existing tags do not meet your requirements, please go to the Console to manage tags.
Source database settings	Source database type	The source database type selected at the time of purchase. It cannot be modified.
	Region	The source database region selected at the time of purchase. It cannot be modified.
	Access type	Please choose according to your scenario. This scenarios takes "TencentDB" as an example. For preparation work of different access types, please refer to Preparation Overview . Public network: The source database can be accessed via a public network IP. Self-Built on CVM: The source database is deployed on Tencent CVM . Direct connect: The source database can be connected to Tencent Cloud Virtual Private Cloud (VPC) via direct connect . VPN access: The source database can be connected to Tencent Cloud VPC via VPN connections . TencentDB: The source database is a TencentDB instance. CCN: The source database can be connected to Tencent Cloud VPC via Cloud Connect Network .
	Cross-account/intra-account	This Account: The source database instance and the target database instance belong to the same Tencent Cloud root account. Cross-account: The source database instance and the target database instance belong to different Tencent Cloud root accounts. The following takes intra-account migration as an example. For cross-account operations, please refer to TencentDB Cross-Account Migration Guide .
	Database instance	Select the instance ID of the source database.
	Account	The account of the source SQL Server database. The account must have the required permissions.

	Password	The password of the source SQL Server database account.
Target database settings	Target database type	The target database type selected at the time of purchase. It cannot be modified.
	Region	The target database region selected at the time of purchase. It cannot be modified.
	Access Type	Select according to your scenario. For this scenario, select "TencentDB".
	Database instance	Select the instance ID of the target database.
	Account	The database account of the target database. It must have the required permissions.
	Password	The password of the target database account.

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

✓ Set source and target databases

2 Set migration options and select migration objects

3 Verify task

Migration Type *

Full migration

Full + Incremental migration

Migration Object *

Specify object

ⓘ The search results of source database objects can display up to 200 records. For more objects, please specify the object name for targeted search.

Source Database Object

Search database name, supporting fuzzy match

Refresh Select All Clear Reset

Selected Object ⓘ

Unfold all Fold all

ⓘ For migration notes, please see [Migration FAQs](#)

Previous

Save

Configuration Item	Description
Migration type	<p>Please choose according to your scenario.</p> <p>Full migration: The entire database is migrated. The data to be migrated includes only the existing content of the source database at the time of task initiation and does not include the incremental data written to the source database after the task initiation.</p> <p>Full + incremental migration: The data to be migrated include the existing content of the source database when the task is initiated as well as the incremental data written to the source database after the task is initiated. If there are data writes to the source database during migration, and you want to smoothly migrate the data in a non-stop manner, select this option.</p>
Migration object	<p>Only database-level migration is supported, meaning that all objects in the specified database need to be migrated together. Select the database to be migrated from the source database objects, then move it to the Selected Object box.</p>

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

If the task verification fails, refer to [Pre-Verification Failure Handling](#), fix the issue, and reinitiate the verification task.

Failed: Indicates that the verification items failed the check, the task is blocked, and you need to fix the problem and run the verification task again.

Alarm: Indicates that the verification items do not fully meet the requirements. You can continue with the task, but it may have some impact on the business. Users need to evaluate based on the prompts whether to ignore the warning or fix the issues before continuing.

Set source and target databases > Set migration options and select migration objects > **3 Verify task**

Task ID / Name	Running Mode	Source Database Type	Target Database Type	Source Access Type	Target Access Type	Address
	Immediate execution	SQL Server	SQL Server	Database	Database	Source: (mtdtsqjany/tset1) Target: (mtdtsqjany/tset2)

Migration Type: Full + Incremental migration
Migration Object: Specify object ①
Unfold all Fold all

test1 (Entire database selected)

Create Verification Task

Query Verification Result

- ✓ Migration parameters are correct
- ✓ Migration network is available
- ✓ Source database is accessible
- ✓ Target database is accessible
- ✓ Disk capacity is sufficient
- ✓ Character sets are consistent
- ✓ Database Version Check
- ✓ User permissions are sufficient
- ✓ Source database exists
- ✓ Target database exists

Previous Verify Again **Start Task** Start Later

6. Return to the data migration task list. The task enters the ready-to-run state. After 1 to 2 minutes, the data migration task will officially start.

If you need to view the task, delete the task, or perform other operations, please click the corresponding task and perform operations in the **Operation** column. For details, refer to [Task Management](#).

If there is an error in the task, please refer to [Error Handling](#).

7. Assess whether to end the task.

Select **Full migration**: Once the task is completed, it will end automatically; no manual action is required.

Select **Full + incremental migration**: After full migration is completed, it will automatically enter the incremental data synchronization stage. Incremental data synchronization will not end automatically; you need to verify the migration results and manually click **Complete** to end incremental data synchronization. If business switching is needed, refer to [Cutover Instructions](#).

Post-migration Operations

After completing the migration using DTS, it is recommended to perform the following checks on the target database: Permission completeness. Permissions will affect operations performed on the database. The migration only restores data. To restore other service-level permissions, such as database users and login user names, you need to recreate them and associate them with database accounts.

Indexes: Reindexing is recommended. As the physical environment of the data files changes, database index statistics may not be updated in a timely manner. It is advised to perform reindexing; otherwise, database performance may degrade.

Instance-level objects. After the migration is completed, users need to re-create these by themselves.

Related APIs

For DTS-related APIs, please refer to: [Viewing Related APIs](#).

Cloud Database Across-Account Migration Guide

Last updated : 2024-08-02 17:28:32

This document describes how to migrate data between instances with Data Transfer Service (DTS) across Tencent accounts.

Application Scope

The source database is an instance of Tencent Cloud's cloud database.

Prerequisite

You have created the target database instance.

Note

This operation involves multiple account information configuration items. The following lists the main configuration logic for easier understanding and configuration.

Data migration direction: Source database (database instance under another account) > target database (database instance under the current account).

The account executing the migration task can be the root account or a sub-account of the target database.

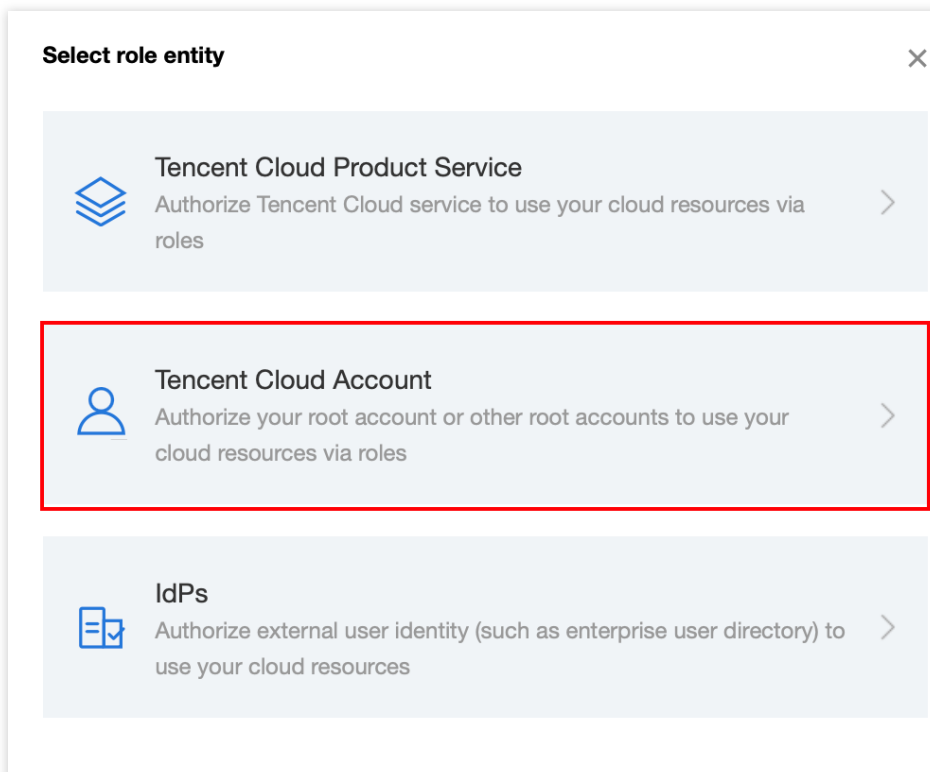
Use the root account to execute the migration task: Before executing the task, ask the root account of the source database to grant the root account of the target database access to the source database.

Use the sub-account to execute the migration task: Before executing the task, ask the root account of the source database to grant the root account of the target database access to the source database by role, and then ask the root account of the target database to grant its sub-account access to the source database.

Authorizing Account

To execute the migration task with a root account or a sub-account, follow steps 1–6 or steps 1–11 respectively.

1. Log in to the [CAM console](#) with the Tencent Cloud root account of the source database. If the sub-account has CAM and role permissions, you can also log in with the sub-account.
2. Click **Roles** on the left sidebar to enter the **Role Management** page. Then, click **Create Role**.
3. On the **Enter Role Entity Info** page, select **Tencent Cloud Account**.



4. On the **Enter Role Entity Info** page, configure the relevant information and click **Next**.

The screenshot displays the "Enter Role Entity Info" page, which is the first step in a four-step process (indicated by the progress bar at the top: 1. Enter Role Entity Info, 2. Configure Role Policy, 3. Set Role Tag, 4. Review).

Fields and options on the page include:

- Tencent Cloud Account Type**: Two radio buttons. ☐ Current root account and ☒ Other root account.
- Account ID**: A text input field containing "2" followed by masked characters and "6", with a green checkmark icon to its right.
- Console access**: A checkbox labeled "Allow the current role to access console".
- External ID**: A checkbox labeled "Enable verification (You're advised to enable this feature when a third-party platform uses this role.)".

A blue information box at the bottom provides details about the External ID:

i The external ID is a string of characters that you define for this role. To use this role, a user needs to pass in this external ID as you set. This improves the security of role assuming by preventing unauthorized use of the role when the role information is leaked or guessed. You're advised to enable external ID verification if you will allow a third-party platform to use the role to be created, or if the account and role informat

A blue **Next** button is located at the bottom left of the form.

Tencent Cloud Account Type: Select **Other root account**.

Account ID: Enter the Tencent Cloud root account ID of the target database, which can be viewed in [Account Information](#). Even if the target database is owned by a sub-account, you still need to enter the root account ID here.

External ID: You can enable it as needed.

Note:

If an external ID is used, keep it properly, as it cannot be queried in DTS.

5. On the **Configure Role Policy** page, select the corresponding policies of DTS and the source database and click **Next**.

Select `QcloudDTSReadOnlyAccess` as the DTS service policy.

For policies corresponding to the source database service, select the read-only service policy and the list acquisition policy of the source database.

If the source database is TencentDB for SQL Server, add `QcloudSQLServerReadOnlyAccess` to obtain its read-only access permission.

Note:

`QcloudCDBReadOnlyAccess` must be added for the source database, otherwise you cannot obtain its instance list information while configuring the migration task.

The screenshot shows the 'Configure Role Policy' step in a four-step process. The 'Select Policies (8 Total)' list on the left includes:

Policy Name	Policy Type
<input type="checkbox"/> Full read-write access to TSI	Preset Policy
<input type="checkbox"/> QcloudAccessForDTSRole This policy is for the DTS service role(DTS_QCSRole) to be associated and used by DTS t...	Preset Policy
<input type="checkbox"/> QcloudDTSReadOnlyAccess QcloudDTSReadOnlyAccess	Preset Policy
<input type="checkbox"/> QcloudTSFFullAccess QcloudTSFFullAccess	Preset Policy
<input checked="" type="checkbox"/> QcloudTSFReadOnlyAccess QcloudTSFReadOnlyAccess	Preset Policy

Support for holding shift key down for multiple selection

On the right, the '2 selected' list shows:

Policy Name
QcloudTSFReadOnlyAccess
QcloudTSFReadOnlyAccess
QcloudSQLServerReadOnlyAccess
Read-only access to TencentDB for SQL Server

Navigation buttons: Back, Next

6. Configure role tags, set the role name on the **Review** page, and click **Complete**.

Note:

Record the configured role name, which needs to be entered when you create the migration task later.

✓ Enter Role Entity Info > ✓ Configure Role Policy > ✓ Set Role Tag > 4 Review

Role Name *

DTS-role

Description

Role Entity

Account-

Access Type

Programming access

Tag

DTSTest:auto_test

Policy Name	Description	Policy Type
QcloudTSFReadOnlyAccess	QcloudTSFReadOnlyAccess	Preset Policy
QcloudSQLServerReadOnlyAcc...	Read-only access to TencentDB for SQL Server	Preset Policy

Back


Complete


Note:


To execute a migration task with the root account, just follow the steps above; to execute a migration task with a sub-account, you need to ask the root account to authorize the sub-account as following steps 7-11.

7. (Optional) Log in to the [CAM console](#) with the Tencent Cloud root account of the target database and click **Policies** on the left sidebar. Then, click **Create Custom Policy** on the right and select **Create by Policy Syntax**.

Select Policy Creation Method

**Create by Policy Generator**
Select service and actions from the list to auto-generate policy syntax

**Create by Policy Syntax**
Write policy syntax to generate policies

**Authorize by Tag**
Grant permissions of resources under certain tags to users or user groups

8. (Optional) Select **Blank Template** and click **Next**.

1 Select Policy Template > **2 Edit Policy**

Template Type: All Templates Search policy name

Select a template type

All Templates (831 Total)

☒ **Blank Template**

☐ **AdministratorAccess**
This policy allows you to manage all users under your account and their permissions, financial information and cloud assets.

☐ **QCloudAutoAccess**
This p (Exc

☐ **ReadOnlyAccess**
This policy authorizes you with the read-only access to all cloud assets that support authentication at API or resource level in your...

☐ **QCloudFinanceFullAccess**
This policy allows you to manage all financial items in your account, such as payment and billing.

☐ **QcloudAutoAccess**
Auto t

☐ **QcloudAccessForCLSRoleInAccessKMS**
This policy is for the CLS service role (CLS_QCSRole) to be associated and used by CLS to access other cloud service resourc...

☐ **QcloudAccessForCLSRoleInCisETL**
This policy allows CLS(Cloud Log Service) DataTransform task to play a role as the clients themselves to Access their Log Topics in...

☐ **QcloudAutoAccess**
This p the cli

Next

9. (Optional) Create a policy and enter the policy name and description as needed. After copying the sample code to the **Policy Content**, replace the content in the red box with the actual information.

✓

Select Policy Template

>

2

Edit Policy

Policy Name *

policygen-20230707104851

After the policy is created, the policy name cannot be modified

Description

Policy Content

[Use Legacy Version](#)

```
1  {
2    "version": "2.0",
3    "statement": [
4      {
5        "effect": "allow",
6        "action": ["name/sts:AssumeRole"],
7        "resource": ["qcs::cam::uin/10*****8:roleName,DTS-role"]
8      }
9    ]
10 }
```


role name

The primary account ID of the source database

Previous

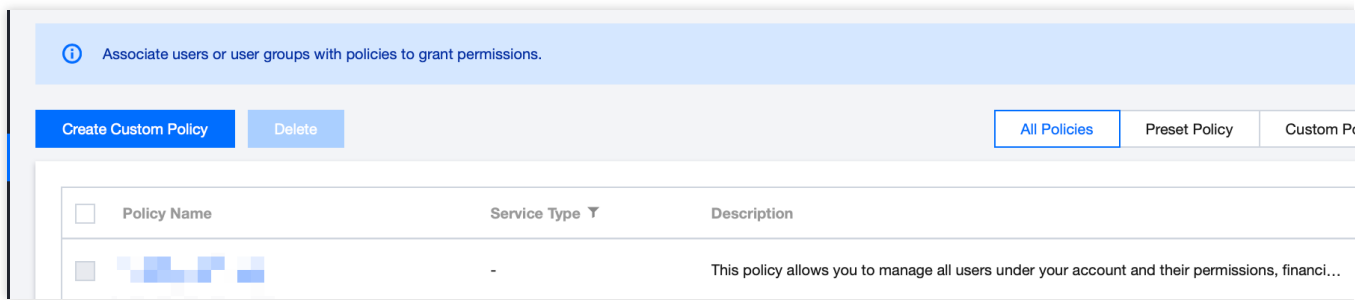
Complete

Sample policy syntax:

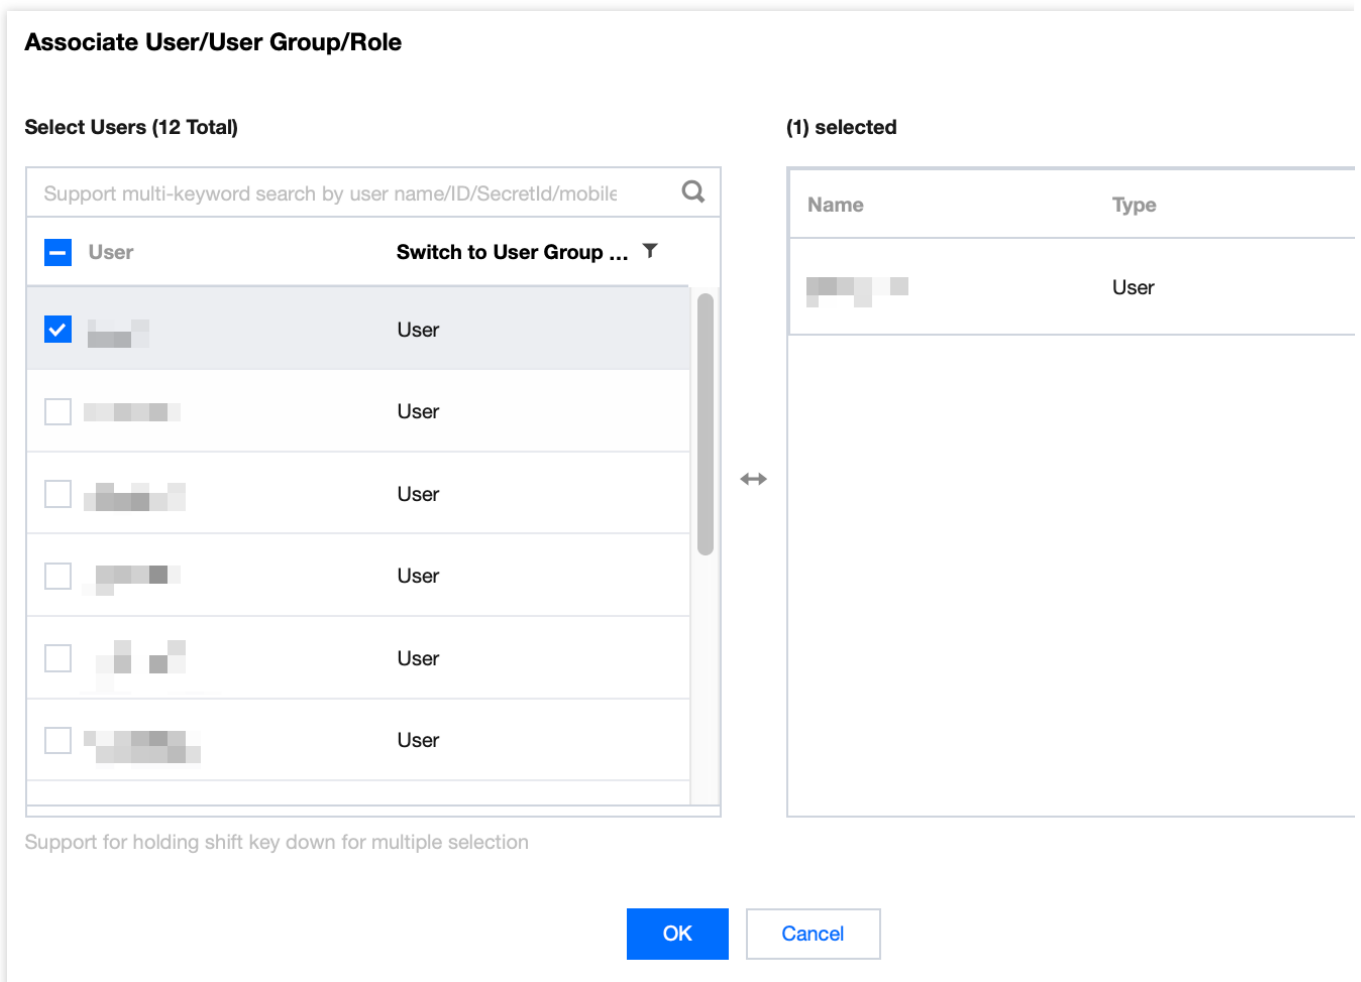


```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "allow",
      "action": ["name/sts:AssumeRole"],
      "resource": ["qcs::cam::uin/10*****8:roleName/DTS-role"]
    }
  ]
}
```

10. (Optional) Click **Complete**, return to the **Policy List** page, and click **Associate Users/Groups**.



11. (Optional) Select the sub-account of the target database instance (which is the sub-account used to execute the migration task) and click **OK**:



Creating Migration Task

1. Log in to the [DTS console](#) with the Tencent Cloud account of the target database instance.
2. Select **Data Migration > Create Migration Task** to purchase a migration task.

3. After making the purchase, return to the data migration task list and click **Configure** in the **Operation** column to enter the migration task configuration page.

4. On the source and target database configuration page, configure the database information.

Configure the key parameters for cross-account data migration as follows:

Access Type: Select **Database**, indicating that the source database is a TencentDB instance.

Cross-/Intra-Account: Select **Cross-account**.

Cross-Account ID: Enter the root account ID of the source database.

Cross-Account Role Name: The **role name** created in step 6 in [Authorizing Account](#). For more information on roles, see [Role Overview](#) and [Cross-account Access Role](#).

Role External ID: This parameter is optional, and its value can be obtained from the previous section. For more information on roles, see [Role Overview](#) and [Cross-account Access Role](#).

Note:

After completing the above configuration, select the **Region** to get the instance list under the source database account. If an error occurs while getting the instances, the configuration may be incorrect, or no authorization has been performed. For more information, see [FAQs](#).

5. On the **Set migration options and select migration objects** page, set the data migration options, select migration objects, and click **Save and Go Next**.

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

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

8. Return to the data migration task list, and the task will be in the "Running" status.

FAQs

1. What should I do if the error "role not exist[InternalError.GetRoleError]" is reported when the instance list is pulled across accounts?

Check whether the **Cross-Account ID** (the root account ID of the source database) and **Cross-Account Role Name** (the **role name** created in step 6 in [Authorizing Account](#)) have been correctly configured.

2. What should I do if the error `InternalError: InternalInnerCommonError` is reported when the database instance list is obtained?

The policy of the Tencent Cloud service to which the source database belongs hasn't been granted to the role. Grant it as instructed in step 5 in [Authorizing Account](#).

3. What should I do if the error "you are not authorized to perform operation (sts:AssumeRole), resource (qcs::cam::uin/1xx5:roleName/xxxx) has no permission" is reported when the instance list is pulled across accounts? Error cause:

The account that you use to create the migration task is a sub-account without the

`sts:AssumeRole` permission.

Solution:

Use the root account to create the migration task.

Ask the root account of the target database to authorize the sub-account as instructed in [Authorizing Account](#) and set `resource` in the policy syntax to the field in blue in the error message.

Pre-Validation Failure Handling

Migration Parameter Check

Last updated : 2024-07-30 17:53:27

Check Details

The target instance cannot contain the database to be migrated; otherwise, an error will be reported with the details:

"The database to be migrated already exists in the target instance."

The provided migration account must be able to be logged in normally; otherwise, an error will be reported with the details: "The provided migration account cannot be logged in."

The name of the database to be migrated cannot contain sensitive characters; otherwise, an error will be reported with the details: "The name of the database to be migrated contains sensitive characters."

The name of the database to be migrated should comply with standards; otherwise, an error will be reported with the details: "The name of the database to be migrated does not comply with standards. Database names must consist of letters, numbers, and underscores, and must start with a letter."

The name of the database to be migrated should be at least 1 character long; otherwise an error will be reported with the error details: "The name of the database to be migrated must be at least 1 character long."

The name of the database to be migrated cannot exceed 64 characters; otherwise an error will be reported with the error details: "The name of the database to be migrated cannot exceed 64 characters."

Fixing Method

Serial Number	Error Details	Fixing Method
1	The database to be migrated already exists in the target instance.	Please change the database to be migrated in the source instance or delete the database with the same name as the database to be migrated in the target instance.
2	The provided migration account cannot be logged in.	<ol style="list-style-type: none">1. Please check if the source instance is already in the Running state.2. Please check if the port of the source instance is being blocked by a firewall security group.3. Please check if the port of the source instance is entered incorrectly.4. Please check if the account and password of the source instance are entered incorrectly.

3	The name of the database to be migrated contains sensitive characters.	Please change the database to be migrated in the source instance or rename the database to remove illegal characters before migration.
4	The name of the database to be migrated does not comply with standards. Database names must consist of letters, numbers, and underscores, and must start with a letter.	Please change the database to be migrated in the source instance or rename the database according to naming standards before migration.
5	The name of the database to be migrated must be at least 1 character long.	Please change the database to be migrated in the source instance or rename the database according to naming standards before migration.
6	The name of the database to be migrated cannot exceed 64 characters.	Please change the database to be migrated in the source instance or rename the database according to naming standards before migration.

Migration Network Check

Last updated : 2024-07-31 09:28:01

Check details

Check if the internal migration network is connected. If the check is not completed, you may see the error message: "Waiting for network check." Please wait for a moment without taking any action. If there is a problem with the internal migration network, you will see the error details: "Network check failed."

Fixing Method

If a migration network check error is reported: "Network check failed," please [submit a ticket](#) to get a solution.

Source Database Connectivity Check

Last updated : 2024-07-31 09:29:33

Check Details

The Data Transfer Service (DTS) server needs to be able to connect to the source database; otherwise an error will be reported with the error details: "Check if the DTS server can connect to the source database."

The original server needs to be connected; otherwise, an error will be reported with the error details: "Insufficient migration account permissions. Unable to complete the migration."

Fixing Method

Serial Number	Error Details	Fixing Method
1	Check if the DTS server can connect to the source database.	Please check if the SQL service startup account for the source instance is using the built-in Local System account. You need to select "Built-in account" in "Log on as" in the startup configuration and modify it to start with the Local System account.
2	Insufficient migration account permissions. Unable to complete the migration.	Please check if xp_cmdshell is enabled on the source instance. xp_cmdshell needs to be enabled. Please ensure the migration account has sysadmin permissions.

Target Database Connectivity Check

Last updated : 2024-07-31 09:31:18

Check Details

The Data Transfer Service (DTS) server needs to be able to connect to the target database; otherwise, an error will be reported with the error details: "Check if the DTS server can connect to the target database."

Fixing Method

If a target database connectivity check error is reported: "Check whether the DTS server can connect to the target database," please [submit a ticket](#) to get a solution.

Disk Space Check

Last updated : 2024-07-31 09:32:19

Check Details

Check whether the disk space on the server of the source database is sufficient. During migration of a self-built source instance, at least 50 GB disk space should be reserved on the server of the source database. We recommend you reserve a space at least 1.5 times the size of the migration data to avoid errors; otherwise, an error will be reported with the error details: "Check whether the disk space on the destination server is sufficient."

Fixing Method

If a disk space check error is reported: "Check whether the disk space on the destination server is sufficient," please check the purchased disk space size of the target instance. It is recommended that the disk space of the target instance is not smaller than the disk space of the source instance, preferably 1.5 times the size of the source instance's disk space. You can expand the disk space size of the target instance through the [Adjusting Instance Specification](#) feature in the console.

Character Set check

Last updated : 2024-07-31 09:33:22

Check Details

The character sets of the source database and target database need to be consistent, otherwise, an error will be reported with the details: "Check if the character sets are consistent."

Fixing Method

If a character set check error is reported: "Check if the character sets are consistent," please check if the character sets at the database level are consistent. If they are not consistent, please adjust the character sets of the source database and the destination database to be consistent.

Database Version Check

Last updated : 2024-07-31 09:34:41

Check Details

The source database version cannot be later than the target database version; otherwise, an error will be reported with the error details: "Check if the database version numbers are consistent."

Fixing Method

If a data version check error is reported: "Check if database version numbers are consistent," please verify the database versions of the source and target instances. A later version instance cannot migrate to an earlier version instance. For example, a source instance of version 2012 cannot migrate to a target instance of version 2008. The target instance's version must be later than or equal to the source instance's version for migration to proceed. You can use the console's [Adjusting the Instance Version](#) feature to upgrade the target instance's version.

User Permission Check

Last updated : 2024-07-31 09:35:38

Check Details

Check if the target instance permissions exist. If not, an error will be reported with the error details: "Check if target instance permissions exist".

Fixing Method

If a user permission check error is reported: "Check if target instance permissions exist", please [submit a ticket](#) to get a solution.

Source Database Existence Check

Last updated : 2024-07-31 09:36:42

Check details

Check if the source database to be migrated no longer exists. If it does not exist, an error will be reported with the error details: "Check if the source database exists".

Fixing Method

If a source database existence check error is reported: "Check whether the source database exists," please confirm whether the database to be migrated exists in the source instance. If it does not exist, you need to reconstruct the migration task. It is recommended not to delete the source database during the migration process.

Target Database Existence Check

Last updated : 2024-07-31 09:37:47

Check Details

When migrating the source database to the target instance, ensure that there are no databases with the same name as the source database to be migrated in the target instance; otherwise, an error will be reported with the error details: "Check if the target database exists."

Fixing Method

If a target database existence check error is reported: "Check if the target database exists," it means that the database to be migrated already exists in the target instance. Please delete or rename the database with the same name in the target instance and then retry.

Error Handling

Last updated : 2024-07-31 09:39:35

Problem Scenario

After the migration task officially starts, it will go through the following **Migration Steps**.

1. Prohibit backup Jobs.
2. Back up databases.
3. Transfer backup files.
4. Restore databases.
5. Deploy real-time synchronization (This step is involved only when the migration type is: full + incremental migration).

During the above migration steps, if a task failure occurs, click **Error Message** to understand the cause of the task failure, and click **Error Details** for handling suggestions.

Problem Handling

The following table lists the error messages of migration task failures and handling suggestions.

Serial Number	Error Message	Suggestions
1	Task failed. Please submit a ticket to get a solution.	Please submit a ticket to get a solution.
2	The database to be migrated was not found in the source instance.	The database to be migrated was not found in the source instance. Please confirm whether the database to be migrated exists in the source instance.
3	The migration initialization operation cannot be performed on the source instance.	The migration initialization operation cannot be performed on the source instance: <ol style="list-style-type: none">1. Please check if the SQL service startup account for the source instance is using the built-in Local System account. You need to select "Built-in account" in "Log on as" in the startup configuration and modify it to start with the Local System account.2. Please check if the source instance has xp_cmdshell enabled. You need to enable xp_cmdshell.

		3. Please ensure that the migration account has sysadmin permissions.
4	Later-version instances cannot migrate to earlier-version instances.	Later-version instances cannot migrate to earlier-version instances, e.g., a source instance of version 2012 cannot migrate to a target instance of version 2008. The target instance version must be greater than or equal to the source instance version for migration to occur. You can upgrade the target instance version using the Adjusting Instance Version feature in the console.
5	Backup jobs for the source instance cannot be disabled.	Please submit a ticket to get a solution.
6	Creating full backup files for the source database failed.	<p>Creating full backup files for the source database failed:</p> <ol style="list-style-type: none"> 1. Please execute net share in the cmd terminal of the source instance to check if a shared folder named backup exists. 2. Please check if the backup files are generated in the shared folder of the source instance. 3. Please check if the source instance does not have enough disk space to create a backup.
7	Source instance failed to transfer backup files.	<p>Source instance failed to transfer backup files:</p> <ol style="list-style-type: none"> 1. Please check if the file-sharing service is enabled on the source instance. 2. Please check if the system account beginning with "ls" is correctly created in the source instance. 3. Please execute net share in the cmd terminal of the source instance to check if a shared folder named "backup" exists. 4. Please check if the system account beginning with "ls" has full control permissions for the shared folders.
8	Failed to recover the database using the backup file on the target instance.	<p>Failed to recover the database using the backup file on the target instance:</p> <ol style="list-style-type: none"> 1. Please check if the file-sharing service is enabled on the source instance. 2. Please check if the system account beginning with "ls" has been correctly created in the source instance. 3. Please execute net share in the cmd terminal of the source instance to check if a shared folder named 'backup' exists. 4. Please check if the system account beginning with "ls" has full control permissions for the shared folders.
9	Incremental synchronization deployment failed.	Please submit a ticket to get a solution.
10	Failed to synchronize	Failed to synchronize incremental logs:

	incremental logs.	<ol style="list-style-type: none">1. Network connectivity exception occurs during synchronization and sharing incremental files failed. Please redo the migration.2. During incremental migration, transaction logs were truncated. Please stop the log backup job on the source instance and redo the migration.
11	After synchronization, data inconsistency exists.	<p>After synchronization, data inconsistency exists. Full backup data synchronization has been completed. Please manually check the migrated data. If there are significant data differences, re-migration is required. The possible causes of data inconsistency are as follows:</p> <ol style="list-style-type: none">1. The source database was not stopped for writes: Before clicking "Complete" for the task, the source database needs to be stopped for writes for 3-5 minutes to avoid data verification inconsistencies.2. Incremental log synchronization failure: During incremental synchronization, log backup and log truncation operations on the source instance need to be stopped.

Data Migration (Legacy)

Migrating Data from CVM-based SQL Server Instance

Last updated : 2024-01-18 17:23:30

Operation Scenarios

TencentDB for SQL Server supports migrating data from a CVM-based self-created SQL Server database to a TencentDB for SQL Server instance. This document describes how to configure and run such a migration task.

Note:

Before migration, please make sure that the SQL Server version of the target instance is not below that of the source instance.

Support the upload of single bak file and tar compressed file for recovery.

The name of the migrated database cannot be the same as that of the TencentDB for SQL Server instance.

Directions

Step 1. Create a migration task

1. Log in to the [TencentDB for SQL Server Console](#) and select **Data Migration (Legacy)** on the left sidebar.
2. Click **Create Task**, enter the task name, source database information, and target database information, and select **CVM-based self-created SQL Server database** as the source instance type.

[< Back](#) | **CreateTasks**

1 Task Initialization

>

2 Sync type and Database Table

Task Configuration

Task Name *

Please enter a task name

The task name can contain up to 60 chars, and cannot include <, >, " and /

Source Database Info

Source instance type *

CDB for SQLServer

Region *

South China (Guangzhou)

SQLServer Instance ID *

Select an instance

Account *

Please enter the account, e.g. root

Password *

.....

Destination Database Information

Destination Database Type *

CDB for SQLServer

Region *

South China (Guangzhou)

Instance ID *

Select an instance

Cancel

Next

3. After clicking **Next**, you need to [configure the source SQL Server instance](#) first and then [configure the migration task](#).

Note:

If the error message "Source instance info checking failed!" is displayed, please check the following items for troubleshooting:

Whether the sa account of the source SQL Server instance exists.

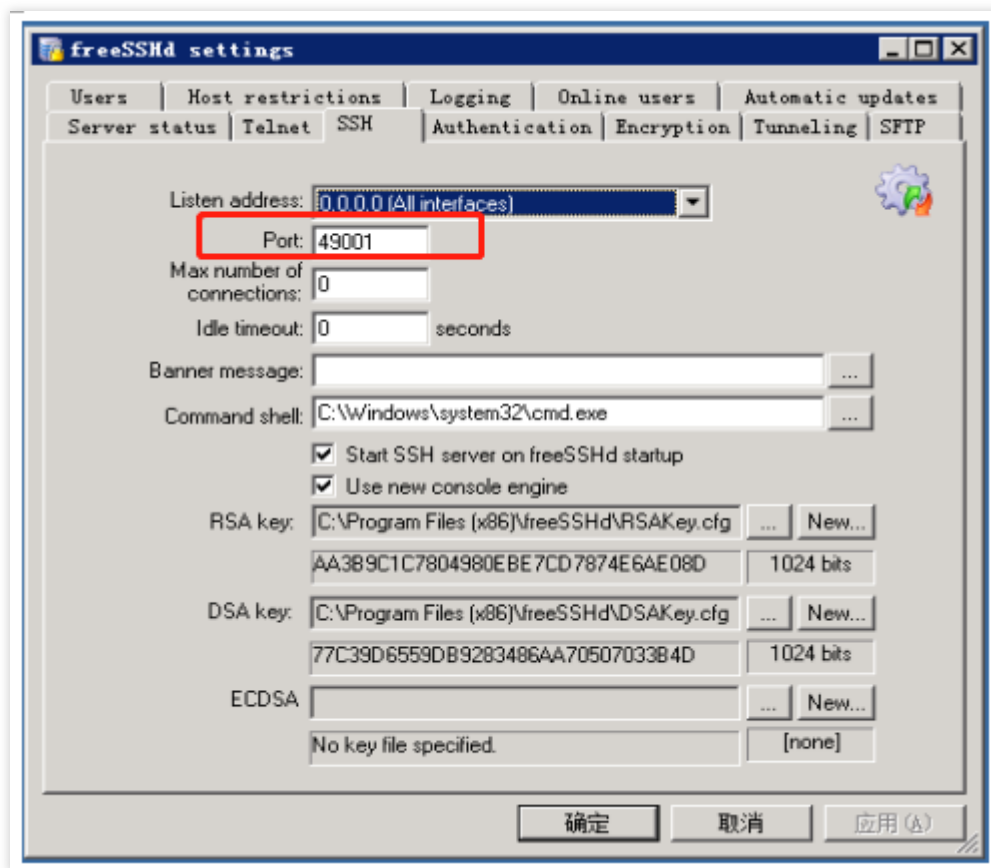
Whether the sa account password of the source SQL Server instance is correct.

Whether the IP and port connectivity of the source SQL Server instance work properly.

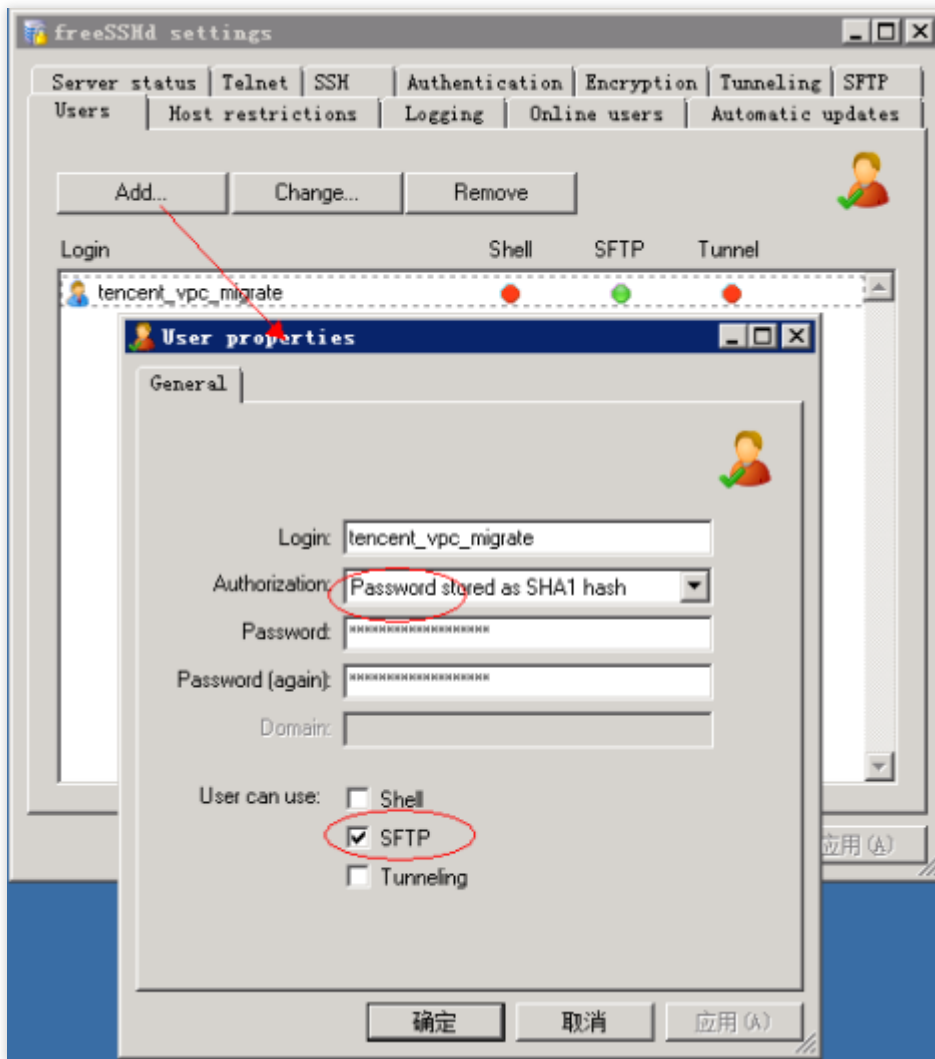
Step 2. Configure the source SQL Server instance

1. Enable the sa account for the source SQL Server instance.

2. Select **Allow remote connections to this server** in **Connections** and set a reasonable remote query timeout period.
3. Select **SQL Server and Windows authentication mode** in **Security**.
4. Enable TCP/IP.
5. Enable the built-in account and select **localsystem**.
6. Allow SQL Server port communication and open port 445 (for the basic network) or 49001 (for a VPC) in Windows Firewall.
7. (Optional) If **VPC** is selected as the **CVM Network**, you need to configure the freeSSHd tool.
8. Download and install freeSSHd with the default options and agree to start the freeSSHd service.
9. Double-click the freeSSHd icon on the desktop. Then, right-click the freeSSHd icon in the taskbar to open the Settings page for configuration.
10. Select the **SSH** tab and set the port to 49001 (the default port here is 22, which should be changed to 49001).



11. Select the **Server status** tab and start the SSH server.
12. Select the **Authentication** tab and select **Allowed** for password authentication.
13. Select the **Users** tab and enter the user `tencent_vpc_migrate` (this username cannot be changed) and the password `tencent_vpc_migrate` (this password cannot be changed) as shown below:



14. Use `D:\\dbbackup\\` (this path cannot be changed) as the backup folder used during SQL Server migration. Select the **SFTP** tab and configure this path as the "SFTP home path".

Step 3. Configure the migration task

Select the migration type, set the database (by selecting the databases/tables to be migrated), and click **Save and Verify**. If the verification fails, you can troubleshoot as prompted.

Step 4. Start the migration task

After the task is created, return to the task list. At this point, the task status is **Initializing**. Select the task and click **Start** to sync the task.

Step 5. Complete the migration task

After the data synchronization is completed (i.e., the progress bar shows 100%), you need to click **Complete** to end the synchronization process. If you selected **Incremental Synchronization** when configuring the migration task, you

need to click **Complete** when the progress bar shows 99%. You can check whether the migration is successful on the **Status**.

If the task status is **task successful*, the data migration is successful.

If the status is **task failed*, the data migration failed. Please check the failure information, fix it accordingly, and then migrate again.

Migrating Data from SQL Server Instance Purchased at Other Cloud Service Provider or Built In-house

Same-Version Migration

Last updated : 2024-01-18 17:23:30

Overview

TencentDB for SQL Server supports data migration using COS files. The migration method described in this document is also applicable to migration from a SQL Server instance purchased at other cloud service provider or built in-house to a TencentDB for SQL Server instance on the same version.

Note:

Before migration, make sure that the SQL Server version of the target instance is the same as that of the source instance.

For the bak files used for migration, make sure that each bak file contains only one database.

The name of the migrated database can't be the same as that of the TencentDB for SQL Server instance.

Full Backup Migration

Preparing a Backup File

There are two ways to prepare a full backup file:

Full Backup After Shutdown

You shut down your SQL Server instance purchased at other cloud service provider or your self-built server, back up the entire database, and then export the backup file (which must be in `.bak` format). The server should be shut down until the migration is completed.

If you choose full backup migration after shutdown, you do not need to perform incremental backup migration separately.

Full Backup Without Shutdown

Note:


The backup file name can't be customized and must follow the naming convention in the script.

After full backup is performed, you need to perform incremental backup restoration until the data on the source instance is the same as that on the target instance.

If the filename contains 1full1, it is full backup.

In the incremental restoration scenario, after backup upload and data migration are completed, the database will be in "Restoring" status before the last incremental backup and restoration is performed. At this time, the database can't be accessed, which is normal. You need to perform the last incremental backup and restoration to make the database accessible.

You do not need to shut down your SQL Server instance purchased at another cloud service provider or your self-built server. Perform backup of the database and export the backup file.



```

---If incremental restoration is not required:
backup database db to disk='d:\\db.bak' with init
---If incremental restoration is required:
declare @dbname varchar(100)
declare @localtime varchar(20)
declare @str varchar(max)
set @dbname='db'
set @localtime =replace(replace(replace(CONVERT(varchar, getdate(), 120 ),'-',''),',','')
set @str='BACKUP DATABASE [' + @dbname + '] TO DISK = N'd:\\dbbak\\' + @dbname +
exec(@str)
go

```

Uploading the Backup File to COS

1. Log in to the [COS console](#).
2. Select **Bucket List** on the left sidebar, and click **Create Bucket**
3. In the pop-up dialog box, configure corresponding information, and click **OK**.

The region of the bucket needs to be the same as that of the SQL Server instance to migrate to.

Cross-region migration with COS is not supported.

4. Return to the bucket list and click the bucket name or **Configuration Management** in the "Operation" column.
5. In **File List**, click **Upload Files**, and you can select one or multiple local files to upload.
6. After the file is uploaded, click the bucket name and obtain the **Object Address** from the basic information in the basic configuration section.

Basic Information

Object Name	
Object Size	
Last Modified	2020-04-28 14:26:17
ETag	
Specified Domain ①	Default Origin Domain ▼
Object Address ①	
Temporary Link ①	Copy Temporary Link Download Objects Refresh

The temporary link carries the signature parameter, and the temporary link can be used to access the object during the validity period of the signature, and the signature is valid for 1 hour (2020-08-25 10:19:57).
Be sure to avoid leaking the temporary link, otherwise your objects may be accessed by other users.

[Migrating Data Via COS File] (id:qianyi_shuju)

1. Log in to the [TencentDB for SQL Server Console](#)
 2. Select **Data Migration (Legacy)** in the left sidebar, and click **Create Task** to create new offline migration task.
- Task Name: Custom.
- Source Instance Type: Select **SQL Server backup recovery (COS mode)**.
- Region: The region of the source database must be the same as that of the source file in COS.
- COS File URL: You can view the file information to get the COS object address after uploading the source file to COS.

Target Database Type and Region: They are automatically generated by the system based on the source database configuration.

Instance ID: Select the target instances that are located in the same region.

Rename Database: Select “Enable or not.”, and you need to enter database name.

Note:

Once enabled, the original database name in the backup file will be reset and then designated as the new database name after it is restored to the cloud database. (Prerequisite: Only one database can be included in the backed up bak file)

Task Configuration

Task Name *

Please enter a task name

The task name can contain up to 60 chars, and cannot include <, >, " and /

Source Database Info

Source instance type *

SQL Server backup recovery (COS mode)

Only COS method is supported. For details, please see ["SQL Server Data Migration by COS"](#)

Region *

South China (Guangzhou)

Original COS file link *

Original COS file link: <http://example-123456.cos>

Enter original COS file link

Destination Database Information

Destination Database Type *

CDB for SQLServer

Region *

South China (Guangzhou)

Instance ID *

Select an Instance

Rename Database *

☒ Enable Database Renaming

When it is enabled, the original database name in the backup file will be reset and the database after being restored to the TencentDB instance will use the new name. (Prerequisite: each .bak file can contain only one database.)

Enter the new database name

Database name is required.

Cancel

Next

3. After completing the configuration, click **Next**.

4. Currently, type and database settings can be changed. Click **Create Task**.

5. Return to task list, and the task status is **Initializing**. Click **Start** above the list to sync the task.

6. After the data sync is completed (i.e., the progress bar shows 100%), you need to click **Complete** at the top of the list to end the sync process. You can check whether the migration is successful on the **Status** column.

If the task status is **Task succeeded*, the data migration is successful.

If the status is **Task failed*, the data migration failed. You need to check the failure information, fix it accordingly, and then migrate again.

Incremental Backup Migration

Preparing a Backup File

Note:

The backup file name cannot be customized and must follow the naming convention in the script.

If the filename contains 1diff1, it is incremental backup.

(Optional) When there are multiple incremental backup files, all of them except the last one can be generated in the following way. They should be uploaded for data migration in sequence; otherwise, the migration will fail.

Note:

After "backup generation, backup upload, and data migration" are completed, the database will be in "restoring" status. At this time, it can't be accessed, which is normal. You can repeat this operation until the last backup upload and migration, and then perform the next step (i.e., "last incremental backup after server shutdown").




```
declare @dbname varchar(100)
declare @localtime varchar(20)
declare @str varchar(max)
set @dbname='db'
set @localtime =replace(replace(replace(CONVERT(varchar, getdate(), 120 ),'-',''),',','')
set @str='BACKUP DATABASE [' + @dbname + '] TO DISK = N'd:\\dbbak\\' + @dbname +
exec(@str)
go
```

If the server is shut down, you need to perform an incremental backup, export the backup files, perform subsequent file upload and data migration.

Note:

Only after this operation is performed can the database be accessed normally.



```
declare @dbname varchar(100)
declare @localtime varchar(20)
declare @str varchar(max)
set @dbname='db'
set @localtime =replace(replace(replace(CONVERT(varchar, getdate(), 120 ),'-',''),',',''),
set @str='BACKUP DATABASE [' + @dbname + '] TO DISK = N'd:\\dbbak\\' + @dbname +
```

```
exec (@str)
go
```

Uploading Backup File and Migrating Data

1. After preparing the backup file, perform subsequent file upload and data migration as instructed in [Uploading Backup File to COS](#) and [Migrating Data Through a Source File in COS](#).
2. After importing the final incremental backup file (i.e., the .bak file containing `_1diff1_1reconvery1`), the status of target instance status will change from read-only to usable, and you can switch your business to the TencentDB for SQL Server instance.

Cross-version Migration

Last updated : 2024-01-18 17:23:30

Overview

TencentDB for SQL Server supports data migration by using COS files. The migration method described in this document is also applicable to migration from an SQL Server instance purchased at another cloud service provider or self-created instance to a TencentDB for SQL Server instance on different versions.

Note:

Before migration, please make sure that the SQL Server version of the target instance is higher than that of the source instance.

For the .bak files used for migration, please make sure that each .bak file contains only one database.

The name of the migrated database cannot be the same as that of the TencentDB for SQL Server instance.

Full Backup Migration

Preparing backup file

There are two ways to prepare a full backup file:

Full backup after shutdown

You shut down your SQL Server instance purchased at another cloud service provider or self-created instance, back up the entire database, and then export the backup file (which must be in `.bak` format). The server should be shut down until the migration is completed.

If you choose full backup migration after shutdown, you do not need to perform incremental backup migration separately.

Full backup without shutdown

Note:

Stop your own data backup and log backup jobs until the migration (full and incremental) is completed.

The backup file name cannot be customized and must follow the naming convention in the script.


After full backup is performed, you need to perform incremental backup restoration until the data on the source instance is the same as that on the target instance.

If the filename contains `2full2`, it is full backup.

In the incremental restoration scenario, after backup upload and data migration are completed, the database will be in "restoring" status before the last incremental backup and restoration is performed. At this time, the database cannot

be accessed, which is normal. You need to perform the last incremental backup and restoration to make the database accessible.

You do not need to shut down your SQL Server instance purchased at another cloud service provider or self-created instance. Perform full backup of the database and export the backup file.



```
declare @dbname varchar(100)
declare @localtime varchar(20)
declare @str varchar(8000)
set @dbname='db'
set @localtime =replace(replace(replace(CONVERT(varchar, getdate(), 120 ),'-',''),',',''),
set @str='BACKUP DATABASE [' + @dbname + '] TO DISK = N'd:\\dbbak\\' + @dbname +
```

```
exec (@str)
go
```

Uploading backup file to COS

1. Log in to the [COS Console](#).
2. Select **Bucket List** on the left sidebar and click **Create Bucket**.
3. In the creation page that pops up, enter the relevant information and click **OK**.
The region of the bucket needs to be the same as that of the SQL Server instance to migrate to.
Cross-region migration with COS is not supported.
4. Return to the bucket list and click the bucket name or **Configuration Management** in the "Operation" column.
5. Select the **File List** page, click **Upload Files**, and select one or more files for upload.
6. After the file is uploaded, click the bucket name and get the **object address** from the basic information in the basic configuration section.

Basic Information

Object Name	
Object Size	
Last Modified	2020-04-28 14:26:17
ETag	
Specified Domain ⓘ	Default Origin Domain ▼
Object Address ⓘ	
Temporary Link ⓘ	Copy Temporary Link Download Objects Refresh

The temporary link carries the signature parameter, and the temporary link can be used to access the object during the validity period of the link. Be sure to avoid leaking the temporary link, otherwise your objects may be accessed by other users.

Migrating data through source file in COS

1. Log in to the [TencentDB for SQL Server Console](#).
2. Select *Data Migration (Legacy)** on the left sidebar and click **Create Task** to create an offline migration task.
Task Name: custom.
Source Instance Type: select **SQL Server backup recovery (COS mode)**.
Region: the region of the source database must be the same as that of the source file in COS.
Link to Source File in COS: you can view the file information to get the COS object address after uploading the source file.
Target Database Type and Region: they are automatically generated by the system based on the source database configuration.

Instance ID: select the target instance. You can only select an instance in the same region.

[Back](#) | **CreateTasks**

1 Task Initialization >


Task Configuration

Task Name * The task name can contain up to 60 chars, and cannot include <, >, " and /

Source Database Info

Source instance type * Only COS method is supported. For details, please see ["SQL Server Data Migration by COS"](#)

Region *

Original COS file link *  Enter original COS file link

Destination Database Information

Destination Database Type *

Region *

Instance ID *

3. After completing the configuration, click **Next**.

4. Currently, type and database settings can be changed. Click **Create Task**.

5. Return to the task list. At this time, the task status is **initializing**. Select the task at the top of the list and click **Start** to sync the task.

6. After the data sync is completed (i.e., the progress bar shows 100%), you need to click **Complete** at the top of the list to end the sync process. You can check whether the migration is successful based on the **status**.

If the task status is **task successful*, the data migration is successful.

If the status is **task failed*, the data migration failed. Please check the failure information, fix it accordingly, and then migrate again.

Incremental Backup Migration

Preparing backup file

Note:


The backup file name cannot be customized and must follow the naming convention in the script.

If the filename contains `2log2` , it is incremental backup.

(Optional) When there are multiple incremental backup files, all of them except the last one can be generated in the following way. They should be uploaded for data migration in sequence; otherwise, the migration will fail.

Note:

After "backup generation, backup upload, and data migration" are completed, the database will be in "restoring" status. At this time, it cannot be accessed, which is normal. You can repeat this operation until the last backup upload and migration, and then perform the next step (i.e., "last incremental backup after server shutdown").



```
declare @dbname varchar(100)
```




```
declare @localtime varchar(20)
declare @str varchar(8000)
set @dbname='db'
set @localtime =replace(replace(replace(CONVERT(varchar, getdate(), 120 ),'-',''),',','')
set @str='BACKUP LOG [' + @dbname + '] TO DISK = N'd:\\dbbak\\' + @dbname + '_' +
exec(@str)
go
```

Shut down the server, perform incremental backup, export the backup file, upload it, and then migrate the data.

Note:

Only after this operation is performed can the database be accessed normally.



```
declare @dbname varchar(100)
declare @localtime varchar(20)
declare @str varchar(8000)
set @dbname='db'
set @localtime =replace(replace(replace(CONVERT(varchar, getdate(), 120 ),'-',''),',','')
set @str='BACKUP LOG [' + @dbname + '] TO DISK = N'd:\\dbbak\\' + @dbname + '_' +
exec(@str)
go
```

Uploading backup file and migrating data

1. After preparing the backup file, perform subsequent file upload and data migration as instructed in [Uploading Backup File to COS](#) and [Migrating Data Through a Source File in COS](#).
2. After the final incremental backup file (i.e., the .bak file containing `_2log2_2reconvery2`) is imported, the status of the target instance will change from read-only to usable, and you can switch your business to the TencentDB for SQL Server instance.