TencentDB for MongoDB

Operation Guide

©2013–2025 Tencent Cloud. All rights reserved.

Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

The complete copyright of this document, including all text, data, images, and other content, is solely and exclusively owned by Tencent Cloud Computing (Beijing) Co., Ltd. ("Tencent Cloud"); Without prior explicit written permission from Tencent Cloud, no entity shall reproduce, modify, use, plagiarize, or disseminate the entire or partial content of this document in any form. Such actions constitute an infringement of Tencent Cloud's copyright, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

Trademark Notice

Tencent Cloud

This trademark and its related service trademarks are owned by Tencent Cloud Computing (Beijing) Co., Ltd. and its affiliated companies("Tencent Cloud"). The trademarks of third parties mentioned in this document are the property of their respective owners under the applicable laws. Without the written permission of Tencent Cloud and the relevant trademark rights owners, no entity shall use, reproduce, modify, disseminate, or copy the trademarks as mentioned above in any way. Any such actions will constitute an infringement of Tencent Cloud's and the relevant owners' trademark rights, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

Service Notice

This document provides an overview of the as-is details of Tencent Cloud's products and services in their entirety or part. The descriptions of certain products and services may be subject to adjustments from time to time.

The commercial contract concluded by you and Tencent Cloud will provide the specific types of Tencent Cloud products and services you purchase and the service standards. Unless otherwise agreed upon by both parties, Tencent Cloud does not make any explicit or implied commitments or warranties regarding the content of this document.

Contact Us

We are committed to providing personalized pre-sales consultation and technical after-sale support. Don't hesitate to contact us at 4009100100 or 95716 for any inquiries or concerns.

Contents

©2013-2025 Tencent Cloud. All rights reserved.

©2013-2025 Tencent Cloud. All rights reserved.

©2013–2025 Tencent Cloud. All rights reserved.

Operation Guide
CAM
CAM Overview

Last updated: 2025-02-08 10:11:29

Cloud Access Management (CAM) helps you securely and conveniently manage access to Tencent Cloud services and resources. With CAM, you can create sub-users, user groups, and roles, and control their access scope through policies. CAM supports SSO capabilities for users and roles, allowing you to set up interoperability between enterprise users and Tencent Cloud based on specific management scenes.
The Tencent Cloud root account you initially created has full access to all services and resources under the account. It is recommended to protect the credentials of the root account, use sub-users or roles for daily access, enable multi-factor authentication, and periodically rotate keys.

Background

If you use multiple cloud platform services, such as Cloud Virtual Machine, Virtual Private Cloud and CloudDB, managed by different people but sharing your cloud account tokens, you might face the issues:

- The risk of your key being compromised is high since multiple users are sharing it.
- You cannot restrict others' access permissions, which can easily lead to accidental operations causing security risks.

Basic Concepts

Root account
When you apply for a Tencent Cloud account, a root account is created by the system which you can use to log in to Tencent Cloud services. A root account is the entity used to bill your usage of Tencent Cloud resources. A root account has full access to all the resources under it by default and can create and authorize sub-accounts.
Sub-account
A sub-account is created by and belongs to the root account. Every sub-account has a definite ID and identity credential.
Identity credentials
Include login credentials and access certificates. **Login credentials** refer to user login names and passwords. **Access certificates** refer to cloud API keys (SecretId and SecretKey).
Resource
A resource is an object manipulated in Tencent Cloud services, such as a TencentDB for MongoDB instance.
Permission
Permissions refer to allowing or denying certain users to perform certain operations. By default, **a root account** has access to all its resources, while **a sub-account** doesn't have access to any resources under the root account.
Policy
A policy is a syntax rule that defines and describes one or more permissions. By default, a sub-account has no access to Tencent Cloud services or resources. To grant a sub-account such access, you need to create a CAM policy.

More information

For more information on access management, see CAM product documentation.

©2013-2025 Tencent Cloud. All rights reserved.

Authorization Policy Syntax

Last updated: 2025-02-08 10:11:51

A policy is a syntax specification of a user permission set that precisely describes the authorized resource set, action set, and authorization conditions.

CAM Policy Syntax

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

The explanation of each statement in the policy syntax is shown in the table below.

Parameter Name	Sub-parameter	Required	Parameter Description
version	No	Yes	Currently, only the value "2.0" is allowed.
statement	effect	Yes	Describes whether the result produced by the statement is "allow" or "explicit denial". It includes two cases: allow and deny.
	action	Yes	Describes the action to be allowed or denied. Operations can be APIs or feature sets (a specific set of APIs described with a permid prefix).
	resource	Yes	This describes the specific data of the authorization. A resource is described in a six-segment format. Detailed resource definitions vary by product.
	condition	Yes	Describes the conditions under which the policy takes effect. Conditions consist of operators, keys, and values. Condition values can include information such as time, IP addresses. Some services allow you to specify additional values in conditions.

Notes:
statement is used to describe the detailed information of one or more permissions. This element includes the permissions or permission sets of multiple other elements such as effect, action, resource, condition, etc. A policy has one and only one statement element.

define action

In CAM policy statements, you can specify any API operation from services that support CAM. For MongoDB, please use APIs prefixed with mongodb:, such as mongodb:BackupDBInstance or mongodb:CreateAccountUser. If you need to specify multiple actions in a single statement, separate them with commas, as shown below:

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

You can also use wildcard characters to specify multiple actions. For example, you can specify all actions whose names begin with the word " Describe " as follows:

©2013-2025 Tencent Cloud. All rights reserved.

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

To specify all MongoDB operations, use the wildcard (*) as shown below:

```
"action": ["mongodb:*"]
```

define resource

Each CAM policy statement has its own resources that apply to it. The general format of resource paths is as follows:

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

- **project_id**: It describes the project information. It is only for compatibility with early CAM logic and does not need to be filled in.
- **service_type**: Short product name, such as mongodb.
- **region**: It indicates regional information, such as bj.
- **account**: The main account information of the resource owner, such as uin/12345678.
- **resource**: Specific resource details of various products, such as instance/instance_id or instance/*.

You can specify the resource in the statement using a specific instance (cmgo-aw6g1g0z) as shown in the following:

```
"resource":[ "qcs::mongodb:bj:uin/12345678:instance/cmgo-aw6g1g0z"]
```

You can also use the wildcard (*) to specify all instances that belong to a specific account as shown in the following:

```
"resource":[ "qcs::mongodb:bj:uin/12345678:instance/*"]
```

If you want to specify all resources, or if a particular API action does not support resource-level permissions, use the wildcard (*) in the resource element, as shown below:

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

If you want to specify multiple resources in one instruction, separate them with commas. The following is an example of specifying two resources:

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

The resources available to MongoDB and their corresponding resource description methods are shown in the table below. Words prefixed with $ are placeholders, where region refers to the region and account refers to the account id.

Resource type	Resource Description Method in Authorization Policies
Instance	`qcs::mongodb:$region:$account:instance/*` `qcs::mongodb:$region:$account:instance/$instanceId`
VPC	`qcs::vpc:$region:$account:vpc/$vpcId`
Security Group	`qcs::cvm:$region:$account:sg/$sgId`

The Default Permission Policy for TencentDB for MongoDB

TencentDB for MongoDB supports the following system permission policies.

Purpose of Policy	Description
`QcloudMongoDBFullAccess`	The administrative privileges for TencentDB for MongoDB grant the sub-account the same permissions as the Tencent Cloud account, including all operation permissions on

©2013-2025 Tencent Cloud. All rights reserved.

	the console and API.
`QcloudMongoDBReadOnlyAccess`	The read-only permission grants the sub-account read-only permission for all resources of the Tencent Cloud account, without operation permissions on the console and API.

The system permission policy `QcloudMongoDFullAccess` is as follows:

```
{
    "version": "2.0",
    "statement": [
        {
            "action": [
                "monitor:GetMonitorData",
                "monitor:DescribeBaseMetrics",
                "mongodb:*"
            ],
            "resource": "*",
            "effect": "allow"
        }
    ]
}
```

The system permission policy `QcloudMongoDBReadOnlyAccess` is as follows:

```
{
    "version": "2.0",
    "statement": [
        {
            "action": [
                "monitor:GetMonitorData",
                "monitor:DescribeBaseMetrics",
                "mongodb:Describe*"
            ],
            "resource": "*",
            "effect": "allow"
        }
    ]
}
```

Custom Permission Policy for TencentDB for MongoDB

TencentDB for MongoDB currently supports the following custom permission policies with resource-level permissions.

> **Notes:**
> Cloud database API operations not listed in the table indicate that the cloud database API operations do not support resource-level permissions. For cloud database API operations that do not support resource-level permissions, you can still grant users permission to use the operation, but the resource element of the policy statement must be specified as *.

Action Name	Description	Resource Description
BackupDBInstance	Backup Database Instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
CreateAccountUser	Creating Account	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId

©2013-2025 Tencent Cloud. All rights reserved.

CreateDBInstance	Creates a monthly subscribed TencentDB instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
CreateDBInstanceHour	Creates a pay-as-you-go TencentDB instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DeleteAccountUser	Delete Account	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeAccountUsers	Query User Information of Account	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeBackupAccess	Get Instance Backup Download Permission	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeBackupRules	Get TencentDB instance backup rule	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeClientConnections	Get number of client connections	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeDBBackups	Query instance backup list	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeDBInstances	Query database instance list	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeInstanceDB	Query database and table information of instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeSlowLog	Get slow log information	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeSlowLogPattern	Get slow log statistics	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
DescribeSpecInfo	Querying available specifications of CloudDB	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
ExchangeInstance	The temporary instance replaces the original instance.	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
IsolateDBInstance	Isolating TencentDB instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
ModifyDBInstanceSpec	Modifying the configuration of TencentDB instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
OfflineIsolatedDBInstance	Deactivating isolated TencentDB instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
RemoveCloneInstance	Deleting a temporary instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
RenameInstance	Renaming instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId

©2013-2025 Tencent Cloud. All rights reserved.

RenewInstance	Renews a TencentDB instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
ResizeOplog	Adjusting instance oplog size	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
RestartInstance	Restart Instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
RestoreDBInstance	Restore database instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetAccountUserPrivilege	Setting user permission	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetAutoRenew	Enabling Auto-Renewal	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetBackupRules	Setting backup rule	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetInstanceFormal	Setting a temporary instance as a formal instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetInstanceMaintenance	Setting instance maintenance window	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetPassword	Set Password	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
SetReadOnlyToNormal	Setting read-only instance as formal instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
TerminateDBInstance	Terminating a monthly subscribed instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
TerminateDBInstanceHour	Terminate a Pay-As-You-Go Instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
UpgradeDBInstance	Upgrading a monthly subscribed instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId
UpgradeDBInstanceHour	Upgrading a pay-as-you-go instance	• qcs::mongodb:$region:$account:instance/* • qcs::mongodb:$region:$account:instance/$instanceId

Example of a Custom Permission Policy

Authorize the account to perform CreateDBInstance and CreateAccountUser operations on the instance cmgo-aw6g****. The example is as follows:

```
{
    "version": "2.0",
    "statement": [
        {
            "effect": "allow",
            "action": [
                "mongodb:CreateDBInstance",
                "mongodb:CreateAccountUser"
```

©2013-2025 Tencent Cloud. All rights reserved.

```
            ],
            "resource": [
                "qcs::mongodb::uin/100001540306:instanceId/cmgo-aw6g****"
            ],
            "condition": {
                "ip_equal": {
                    "qcs:ip": [
                        "10.0.0.4"
                    ]
                }
            }
        }
    ]
}
```

Creating a Custom Permission Policy

For specific operations, please log in to the Cloud Access Management (CAM) console's Policies page and refer to Creating Custom Policies.

©2013–2025 Tencent Cloud. All rights reserved.

Authorization Permission Policy

Last updated: 2025-02-08 10:12:17

Permissions of Tencent Cloud root accounts and sub-accounts are separated. You can grant sub-accounts different permissions as needed, which avoids security risks caused by exposure of your Tencent Cloud account key.

Granting Sub-account permission policy

Background

Company A activates the TencentDB for MongoDB service and wants its team members to manipulate the involved resources. For security or trust considerations, it doesn't want to directly disclose its Tencent Cloud account key to the team members; instead, it wants to create corresponding sub-accounts for them, which can manipulate resources only with authorization by its root account and don't require separate usage calculation and billing, as all fees are charged to its Tencent Cloud account. It also wants to be able to revoke or delete the operation permissions of sub-accounts at any time.

Operation Steps

Step 1. Create a sub-account user

You can create a sub-account user through the console or an API.

- Log in to the Tencent Cloud Access Management (CAM) console and enter the User List page to create a user. For detailed directions, see Creating Sub-user.
- Use access keys to call the AddUser API to add a sub-user and set permissions. For more information, see Adding Sub-user.

(Optional) Step 2. Create a custom permission policy

1. On the Policies page in the CAM console, search for a target policy by policy name in the search box in the top-right corner.
2. If the permission policy does not exist, you need to customize one. For detailed directions, see Creating a Custom Policy.

Step 3. Assign the permission policy to the sub-account

- On the Policies page in the CAM console, find the target permission policy and associate it with the target sub-account. For detailed directions, see Authorization Management.
- On the User List page in the CAM console, find the target sub-account and associate them with the target policy. For detailed directions, see Authorization Management.

More References

Logging in to the Console

You can let your team members use a sub-account to log in to the Tencent Cloud console and access TencentDB for MongoDB. For detailed directions, see Logging in to Console with Sub-account.

Modifying sub-account user information

You can view and modify the information of a sub-account as instructed in User Information.

Deleting sub-account

You can revoke or delete the operation permissions of a sub-account as instructed in Deleting Sub-Users.

Granting Permission Policy Across Tencent Cloud Accounts

Background

Company A activates TencentDB for MongoDB and wants company B to have part of the permissions of its TencentDB for MongoDB operations, such as Tencent Cloud Observability Platform, instance read/write, and slow query operation. Company B wants to have a sub-account to take care of such businesses. In this case, company A can authorize the root account of company B to access TencentDB for MongoDB resources through a role. For the specific concept and use cases of role, see Role Overview.

Operation Steps

©2013-2025 Tencent Cloud. All rights reserved.

Step 1: Company A creating a role for company B

1. Log in to the Tencent Cloud Access Management (CAM) console and go to the Roles page.
2. Click **Create Role**, in the **Select Role Entity** dialog box, select **Tencent Cloud Account**.
3. On the **Create Custom Role** configuration wizard page, create a role.
   a. On the **Enter Role Entity Information** page, select **Cloud Account Type** as **Other Master Account**, enter the primary account of enterprise b in the **Account ID** field, set other parameters as prompted, and click **Next**.
   b. On the **Configure Role Policy** page, select the policy to authorize the role, and click **Next**.
   c. On the **Review** page, in the **Role Name** input box, set the role name, such as DevOpsRole. Review the selected policy, and click **Done**.

Step 2. Company B grants a sub-account the permission to assume the role

1. On the Policies page in the CAM console, click **Create Custom Policy**.
2. In the **Select Policy Creation Method** dialog box, select **Create by Policy Syntax**.
3. In the **Create by Policy Syntax** configuration wizard, create a policy.
   a. In the **Select Template Type** section, select **Blank Template**, and click **Next**.
   b. On the **Edit Policy** page, in the **Policy Name** input box, set the policy name, such as sts:AssumeRole.
   c. In the **Policy Content** section, set the policy content according to the policy syntax, and click **Done**. An example is as follows:

```
{
 "version": "2.0",
 "statement": [
 {
     "effect": "allow",
     "action": ["name/sts:AssumeRole"],
     "resource": ["qcs::cam::uin/12345:RoleName/DevOpsRole"]
 }
 ]
}
```

4. Return to the Policy page, find the created custom policy, and click **Bind User/Group** in the **Operation** column.
5. Associate the custom policy with the sub-account of company B and click **Yes**.

Step 3: Company B using the sub-account to access cloud resources through the role

1. Log in to the console with the sub-account of company B and select **Switch Role** in the profile photo drop-down list.
2. On the role switch page, enter the root account of company A and role name to switch to the role of company A.

More References

- If you need to modify a role, see Modifying a Role.
- If you need to delete a role, see Deleting a Role.
- For more information on how to use CAM, see User Guide.

©2013-2025 Tencent Cloud. All rights reserved.

Managing an Instance
Viewing Instance Details

Last updated: 2025-02-08 10:12:57

Overview

After purchasing TencentDB for MongoDB, you can quickly and intuitively view instance details in the console, including instance running status, capacity usage, cluster master-slave relationship, network status, and more. You can also perform efficient operation and maintenance management on the instance.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The instance is not destroyed and is isolated in the recycle bin. For specific information, see Recycle Bin.

View instance list

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.
5. View the running status, configuration specifications, storage engine, and other information of the target instance.

Instance List Parameters	<0>Parameter meaning<0>
Instance ID/Name	• Instance ID: The unique identification ID of the instance. • Name: The name set when creating the instance. Hover over the instance name and click to re-edit the instance name for easy identification and management.
Monitoring/Status	• Monitoring: Click to quickly view the monitoring metrics data of the instance in the monitoring panel. For specific information, see Viewing Monitoring Data. • Status: The running status of the instance, with the normal status being running. When a task is executing, the currently executing task name will be displayed here, such as configuration changes.

Configuration/ Network	• Specification types: Supports high IO 10G type and cloud disk edition. For specific information, see Product Specification Type.
	• Specifications: configuration specifications of a single node of the instance.
	  ○ Replica set: memory capacity/disk capacity.
	  ○ Sharding instance: memory capacity/disk capacity x shard number.
	• Network: network information of the instance.

Version and Engine	• Database version information: includes 7.0, 6.0, 5.0, 4.4, 4.2, 4.0, 3.6, 3.2. Version 3.2 is no longer available for sale.

©2013-2025 Tencent Cloud. All rights reserved.

	• Storage engine: default is WiredTiger.
Intranet Address	Internal IPv4 address, IP address and port of all Mongod primary/secondary nodes of the database instance. Database instances only support internal network access. When accessing the database through MongoDB Shell, you need to configure the internal IP address and port information. For specific operations, see Connect Instances.
Billing Mode	Billing mode: pay-as-you-go or annual and monthly subscription. The billing methods are different. For details, see Billing Overview.
Used/Total Amount	Disk capacity used by the instance/total disk capacity, making it easy to quickly find the disk usage ratio of the current instance.
Oplog/Shards Information	Click **View/Adjust** to see the reserved disk capacity for Oplog and adjust it according to business needs. For specific operations, see Adjusting Oplog Capacity.
Associated project	The project to which the instance belongs allows you to view all associated instance information. To transfer to another project, see Adjust the project of an instance.
Opertation	• Select **Adjust Configuration** > **Adjust Configuration** to adjust the instance's memory and disk capacity. For specific operations, see Adjusting instance configuration. • Select **Adjust Configuration** > **Node Management** to manage the instance's Mongod and Mongos nodes. For specific operations, see Viewing node information. • Select **More** > **Security Group** to reselect the security group inbound rule. • Select **More** > **Restart** to restart the instance. For specific operations, see Restart Instance. • Annual and monthly subscription billing mode, select **More** > **Return and Refund**; for on-demand instance billing, select **More** > **Destroy** to refund the instance and isolate the instance in the recycle bin. For specific operations, see Terminating Instance. • For on-demand instances, select **More** > **Pay-as-You-Go to Monthly Subscription** to change the billing mode. For details, see Transition from Pay-as-You-Go to Monthly Subscription. • Select **More** > **Management** to view instance details. • Select **More** > **Edit Tags** to modify the instance's tag key-value pairs. For details, see Editing Instance Tags.

Viewing Instance Details

In the **Instance ID/Name** column of the target instance, click the instance ID to enter the **Instance Details** page.

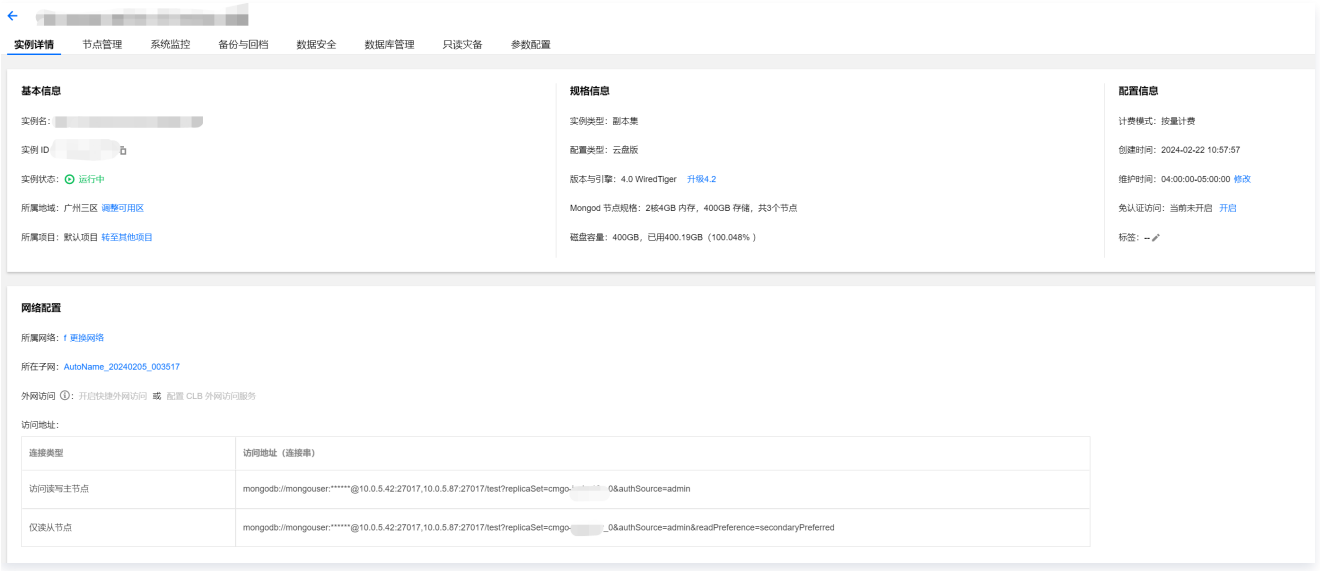

Interface Area	Interface Parameter	Parameter interpretation
Basic Info	**Instance Name**	Custom instance name.
	Instance ID	Unique identification ID of the instance.

©2013-2025 Tencent Cloud. All rights reserved.

	Instance Status	The current running status of the instance, normally: running.
	Region	The region and AZ to which the instance belongs. Click **Modify AZ** to switch to another AZ in the same region. For precautions and specific operations of switching AZ, refer to Modifying Instance AZ.
	Associated project	The project name to which the instance is added. Click **Transfer to Another Project** to reassign the instance to another project. For specific operations, refer to Adjust the Project of an Instance.
Specs Info	Instance Type	The instance cluster architecture type is: replica set or sharded cluster. For detailed information on cluster deployment architecture, refer to System Architecture.
	Configuration Type	Product specification types include: high io 10g type, cloud disk edition. For detailed information, refer to Specification Type.
	Version and Engine	Version information and storage engine of the instance, upgradeable versions. For specific operations, see Version Upgrade.
	Mongod Node Specifications	Configuration information of a single Mongod node, including the number of CPU cores, memory, disk size, and the number of nodes. For details on the specifications supported by replica sets and sharded clusters, see Product Specifications.
	Mongos Node Specifications	Configuration information of a single Mongos node, including the number of CPU cores, memory, and the number of nodes. For details on the specifications supported by replica sets and sharded clusters, see Product Specifications.
	Disk Capacity	Total disk capacity of the instance, used capacity, and disk usage.
Configuration Info	Billing Mode	Billing mode of the instance: pay-as-you-go and annual and monthly subscription.
	Creation Time	Instance creation time.
	Maintenance Time	The instance's maintenance window. To ensure database stability, the back-end system will periodically perform maintenance operations on the instance during the maintenance window. Click **Modify** to adjust the maintenance window. It is recommended to set it during off-peak hours. For details, see Set the instance maintenance time.
	Authentication-Free Access	You can check whether authentication-free access to the database is enabled. If it is currently disabled, click **Enable** to enable authentication-free quick access to the database. For details, see Accessing Instance Without Authentication.
	Tag	Tags associated with the instance. You can modify the tags. For details, see Editing Instance Tags.
Network Configuration	Network	The name of the instance's VPC. Click **switch network** to switch the VPC and subnet. For details, see Switching Instance Network. To create a VPC, see Create a VPC.
	Subnet	The subnet of the instance's VPC. A subnet is specific to an AZ. A VPC allows for subnets in different AZs, and by default, these subnets can communicate with each other via a private network. After modifying the AZ, it is recommended to switch the subnet simultaneously to reduce latency.
	Connection Type	Type of node accessing the database. • Access the primary node: Access the database through the primary node of the instance, which can read and write data. • Access read-only nodes only: Access the database only through read-only nodes. If no read-only nodes are configured when creating an instance, this option will not be displayed. • Access from secondary nodes only: Access the database only through replica nodes. • Access Secondary and Read-Only Nodes Only: Prefer accessing the database through secondary nodes. If secondary nodes fail, access the database through read-only nodes.

©2013-2025 Tencent Cloud. All rights reserved.

	Access Address (Connection String)	Each connection type corresponds to a URI concatenate connection string. You can directly copy the connection string to access the database connect to instance.

More Operations

Renaming the Instance

1. In the Instance List, hover the mouse over the instance name to be modified and click the on its right.
2. Reconfigure an identifiable and manageable instance name in the instance name input box. The requirements are as follows:
   - Character length requirement: [1,60]
   - Chinese characters, English letters, digits, underscores "_", and dashes "–" are allowed.
   - One Chinese character counts as 3 characters, while one English letter, digit, or special symbol counts as one character.

Setting the Instance List Fields

1. Click in the top right corner of the instance list.
2. On the **Customizing List Fields** page, select the fields to display.
3. Click **Yes** to see the newly set fields directly in the instance list.

Instance list export

Click in the top right corner of the instance list to export the entire instance list.

Relevant API

API Name	API Feature Description
DescribeDBInstances	Querying TencentDB instance list
RenameInstance	Modifying Instance Name

©2013–2025 Tencent Cloud. All rights reserved.

Adjusting instance configuration

Last updated: 2025-02-08 10:13:30

Overview

If the configuration of your purchased TencentDB for MongoDB instance isn't suitable for (either below or above) your current business requirements, you can quickly adjust its specifications based on your actual business conditions (at the initial stage, at the rapid development stage, during peak hours, or during off-peak hours), so as to better meet your needs such as full utilization of resources and real-time cost optimization.

Adjusting the configuration specifications

- The memory and CPU cores of mongod and mongos nodes are in fixed combinations, and the disk capacity has a corresponding value range. For example, if the specification of a mongod node is 2-core 4 GB MEM, then the disk capacity can range from 100 to 500 GB.
- Mongod replica nodes: You can select three (one-primary-two-replica), five (one-primary-four-replica), or seven (one-primary-six-replica) nodes in total. Currently, you cannot customize the number of replicas.
- The numbers of mongos nodes supported by single-AZ deployed and multi-AZ deployed instances are different. A single-AZ deployed instance can contain 3-32 nodes, while a multi-AZ deployed instance can contain 6-32 nodes.

For detailed information, please first refer to Product Specifications to help you plan specifications that fit your business. Then, select the required adjustment category from the table below and choose the corresponding operation.

Specification Adjustment Type	Specification Description
Adjust the mongod node specification	You can adjust the memory, disk capacity, and oplog capacity of a mongod node.
Adjust the replica node quantity	You can add or remove replica nodes in both replica set and sharded cluster. You can select three (one-primary-two-replica), five (one-primary-four-replica), or seven (one-primary-six-replica) nodes in total. Currently, you cannot customize the number of replicas.
Adjusting Shard Quantity	A sharded cluster supports increasing the number of mongod shards.
Adjust the mongos node specification	You can adjust the CPU cores and memory of mongos nodes in a sharded cluster.
Adjust the mongos node quantity	You can add mongos nodes in a sharded cluster.
Add read-only nodes	You can add 0-5 read-only nodes.

©2013-2025 Tencent Cloud. All rights reserved.

Accessing Instance Without Authentication

Last updated: 2025-02-08 10:18:38

Background

Auth-free access allows you to access your TencentDB instance quickly and efficiently. However, it also involves security risks. We recommend you enable this feature in test or maintenance scenarios and disable it during business operations.

Version Description

Instances on database versions 3.6, 4.0, 4.2, 4.4, 5.0, and 6.0 support auth-free access.

Notes

- Upgrade to support auth-free access involves kernel upgrade and momentary disconnections.
- Enabling auth-free access will **restart** the instance. Please perform this operation during off-peak hours.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The TencentDB for MongoDB instance is in **Running** status.
- On the instance details page, **Auth-Free Access** is **currently not enabled**.

Enabling auth-free access

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click **Instance ID** to enter the **instance details** page.
6. On the **Instance Details** page, click **Enable** after **Auth-Free Access**.
7. In the **Enable Auth-Free Access** pop-up window, understand the impact after enabling, and click **OK**.

8. On the **Instance Details** page, wait for the instance status to become **Running**. Then, you can connect to the database at the configured private IPv4 address and port without entering the username and password.

Disabling auth-free access

On the **Instance Details** page, click **Disable** after **Auth-Free Access** to disable this feature.

Related Documentation

You can access TencentDB for MongoDB databases by using MongoDB shell or drivers in various programming languages as instructed in Connecting to TencentDB for MongoDB Instance.

©2013-2025 Tencent Cloud. All rights reserved.

Modifying Instance AZ

Last updated: 2025-02-08 10:19:39

Overview

During routine maintenance, you can adjust the MongoDB instance and Cloud Virtual Machine (CVM) to be in the same AZ for lower network latency. TencentDB for MongoDB supports free switching of AZs for cross-AZ instances and also supports adjusting non-cross-AZ instances to cross-AZ instances.

Billing Instructions

Adjusting the AZ will not affect instance billing, please use with confidence.

Notes

Adjusting the AZ will cause a master-slave switch, resulting in a disconnection of about 10 seconds. Please operate during business off-peak period.

Must-Knows

After adjusting the AZ, all attributes, specifications, and connection IP address of the instance will not change. However, switching network will cause the internal network IP of the database to change, requiring application reconnection.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The current instance status is **Running**.
- The target AZ to switch to and the current instance's AZ are in the same region.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the instance ID to enter the **Instance Details** page.
6. On the **Instance Details** page, click **Region** and then **Modify AZ**.

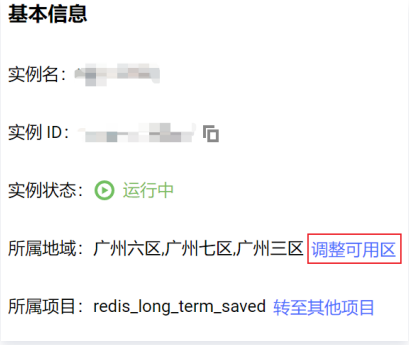

7. In the **Modify AZ** window, understand the impact of modifying the AZ. For multi-AZ deployment instances, configure different AZs for the primary and secondary nodes. For single-AZ deployment instances, configure the same AZ for both primary and secondary nodes.

©2013-2025 Tencent Cloud. All rights reserved.

Note:
Modifying the AZ will cause a master-slave switch, resulting in a disconnection of about 10 seconds. Please perform this operation during off-peak hours.

8. Select the time to execute the AZ switch task after **Switch Time**. Click **understand switch time** to change the instance maintenance window according to the business off-peak period. For specific operations, see adjust instance maintenance time.

Note:
If you have scheduled a task to adjust AZs during the maintenance timeframe, avoid selecting **upon completion of adjustment** to execute immediately before the maintenance time, as it may cause program errors. Initiated tasks cannot be stopped manually; you need to Submit a Ticket to apply.
- **Upon completion of adjustment**: Execute immediately after configuration is completed.
- **Maintenance timeframe**: Execute tasks within the maintenance period.

9. Confirm the **total cost**, click **Yes** to enter the product order page, verify the order, click **Submit Order**, pay the fee, and complete the operation.
The instance status changes to **Switching AZs**. Wait for the task to complete to see the adjusted AZ.

Subsequent Operations

After adjusting AZs, switch the VPC subnet to avoid high access latency. For specific operations, refer to Switching Instance Network.

©2013-2025 Tencent Cloud. All rights reserved.

Set the instance maintenance time

Last updated: 2025-02-08 10:20:21

TencentDB for MongoDB allows you to adjust the instance maintenance time in the console to meet the changes in your business requirements.

Overview

Maintenance time is a very important concept for TencentDB for MongoDB. To ensure the stability of your TencentDB for MongoDB instance, the backend system performs maintenance operations on the instance during the maintenance time. We highly 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.
In addition, we also recommend you perform operations involving data migration, such as adjusting the memory specification or AZ of mongod or mongos nodes, during the maintenance time. Taking the database instance AZ change as an example, as the full and incremental data needs to be synced from the original AZ to the new AZ, data migration will be involved. After the AZ is changed, a momentary disconnection from the database may occur. When the task is initiated, the **Switch Time** can be selected as **During maintenance time**, so that the instance data migration will be started during the next **maintenance time** after the data sync is completed.

> **Notes:**
> Before maintenance is carried out for TencentDB for MongoDB during the maintenance time, notifications will be sent to the contacts configured in your Tencent Cloud account through SMS and email.

Version Description

Currently, TencentDB for MongoDB 7.0, 6.0, 5.0, 4.4, 4.2, 4.0, 3.6, and 3.2 support maintenance time configuration.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the ID of the target instance to enter the **Instance Details** page.
6. On the **Instance Details** page, click **Modify** on the right of **Maintenance Time**.
7. In the **Modify Maintenance Time** dialog box, set the **Start Time** and **Duration**.

8. Click **OK** to complete the operation. You can view the newly set maintenance time on the **Instance Details** page.

©2013-2025 Tencent Cloud. All rights reserved.

Adjust the project of an instance

Last updated: 2025-02-08 10:20:49

TencentDB for MongoDB allows you to assign an instance to a new project in the console to meet your requirements in ever-changing business scenarios.

Background

A project is a set of applications or services that share resources. Applications, services, and resources in different projects are isolated from and don't affect each other, and each project is unique.
You can specify an appropriate project for your database instances to facilitate collaboration. In this way, you can easily manage your instances globally and stay on top of the operational conditions of the entire project.

Version Description

Currently, TencentDB for MongoDB 7.0, 6.0, 5.0, 4.4, 4.2, 4.0, 3.6, and 3.2 support instance project adjustment.

Description

Updating the database project will not incur additional costs.

Must-Knows

Assigning and reassigning TencentDB instances among projects will not affect the services provided by the instances.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- You have specified a project for the instance. The **Default Project** is used by default.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. In the **Basic Information** section, click **Switch to Another Project** on the right of **Project**.
7. On the **Assign to Project** page, select the target project for the instance.
8. Click **OK**. In the **Basic Information** section, **Instance Status** will be displayed on the right of 正在转移项目.
9. Wait for the instance transfer project to be completed. On the right of **Project**, you can see the reassigned project.
   You can filter instances by **Project** in the instance list to view the running status of each instance in the entire project.

API

API Name	API Feature Description
AssignProject	Specify the project to which the cloud database instance belongs.

©2013-2025 Tencent Cloud. All rights reserved.

Editing Instance Tag

Last updated: 2025-02-08 10:21:13

TencentDB for MongoDB supports editing instance tags in the console, making it easier to manage instances through tags.

Background

Tags consist of tag keys and tag values, which can mark TencentDB for MongoDB instances. If you have various cloud resources under your Tencent Cloud account, with multiple associations between different resource types, and the number of cloud resources is increasing, the management complexity will change accordingly. You can use tags to group and categorize resources with the same or related functions. During daily operation and maintenance or when locating problems, you can quickly search for resources based on tags, perform batch operations, and achieve efficient operations and maintenance.

Version Description

TencentDB for MongoDB versions 7.0, 6.0, 5.0, 4.4, 4.2, 4.0, 3.6, and 3.2 all support tag management.

Description

Tag management is one of the free services provided by the cloud platform for your cloud platform account at no additional charge. You can directly access the Control Console to use the product.

Must-Knows

- One tag contains one tag key and one tag value (tagKey: tagValue).
- A maximum of 50 tags can be bound to one instance.
- For each instance, a tag key can correspond to only one tag value.

Prerequisites

You have created a TencentDB for MongoDB instance.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Choose any of the following methods to enter the **Edit Tag** page.
   - In the **Operation** column of the target instance, select **More** > **Edit Tag**.
   - Click the target instance ID, in the **Instance Details** page, under the **Configuration Information** section, click **Tag** on the right side of  .

©2013-2025 Tencent Cloud. All rights reserved.

6. In the **Edit Tag** page, select the appropriate tag key in the **Tag Key** dropdown list and select the corresponding tag value in the **Tag Value** input box.

7. (Optional) If an existing tag does not meet your business requirements, perform the following:
   7.1 On the top right of the current page, click on **Manage tags**.
   7.2 On the **Tag List** page, click on **Create Tag**.
   7.3 On the **Create Tag** page, understand the precautions for setting tags.
   7.4 In the **Tag Key** input box, set a new tag value. In the **Tag Value** input box, input the corresponding tag value. Tag key setting requirements are as follows:
   - Character length requirements [1,63].
   - You can input English letters, digits, and Chinese characters.
   - You can input special symbols: plus sign "+", equal sign "=", underline "_", dashed line "-", English dot ".", English colon ":", slash "/", at "@", English brackets "()", Chinese brackets "()", square brackets "[]", and Chinese square brackets "【】".

   7.5 Click **OK** to complete creation.
   7.6 Then, return to the **Edit Tag** page of the database instance. In the dropdown list of tag keys, click on **reload** to select the newly created tag key and then select the corresponding tag value.
8. Click on **OK** to complete the settings.

More References

For more information on tag management, see Tag Management.

©2013-2025 Tencent Cloud. All rights reserved.

Restarting an Instance

Last updated: 2025-02-08 10:21:35

When the instance encounters connection number full or performance issues, you need to manually restart the instance. This article introduces the specific operations for restarting replica set and sharded cluster instances.

Background

When the system is completely unavailable due to high load, you can restore node operation by using the restart feature. Determined by the architecture of TencentDB for MongoDB, restarting a MongoDB instance is divided into restarting mongos and restarting mongod.

- **mongos**: a routing service configured for MongoDB sharding, which handles query requests from the application layer and determines the location of data in the sharded cluster.
- **mongod**: the primary daemon process for the MongoDB system, which handles data requests, manages data access, and performs background management operations.

Version Description

- Currently, MongoDB 7.0, 6.0, 5.0, 4.4, 4.2, 4.0, 3.6, and 3.2 support instance restart.
- The replica set **4.0** version has a simplified architecture, removed the mongos component, and restarting the instance does not involve restarting the mongos component.

Notes

- During the restart period, 1–2 intermittent disconnections may occur, each lasting about 10 seconds. It is recommended that the program has an automatic reconnection feature.
- During the restart period, this restart operation is unable to cancel. Please operate with caution.

> **Notes:**
> Restarting the entire instance's mongod process carries certain risks. Please contact the MongoDB team for assistance, or restart by node. Refer to Node Restart.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB instance is in **Running** status.

Operation Steps

Restarting a Single Instance

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the row of the target instance, click **Operation** column's **More** > **Restart**.

6. In the **Restart MongoDB** window, click **Viewing Details** to confirm the information of the instances to be restarted.
7. Select the components that need to be restarted, and click **Yes**.
8. In the instance list, you can see the instance entering **Restarting**. Wait for the task to complete.

Batch Restarting Instances

1. In the instance list, select the instances that need to be restarted.

©2013–2025 Tencent Cloud. All rights reserved.

2. Above the entire instance list, click **Restart**.
3. In the **Restart MongoDB** window, click **Viewing Details** to confirm the information of all instances to be restarted.
4. Select the components that need to be restarted, click **Yes**, and wait for the task to complete.

Terminating Instance

©2013–2025 Tencent Cloud. All rights reserved.

Last updated: 2025-02-08 10:22:00

If you no longer need a TencentDB for MongoDB instance, you can directly terminate and return it in the console.

Background

You can terminate an instance if you no longer need it. The terminated instance will be put into the recycle bin. For instances in the recycle bin, you can renew, restore, or release them as needed based on different scenarios.

Version Description

Currently, TencentDB for MongoDB 7.0, 6.0, 5.0, 4.4, 4.2, 4.0, 3.6, and 3.2 support instance termination.

Billing Instructions

- After self-service return, once the instance status changes to **Isolated**, no relevant fees will be incurred.
- The five-day unconditional self-service refund is returned to the original cloud platform account.
- The common self-service refund is returned to your cloud platform account based on the proportion of cash and free credits used for the purchase.
- After a monthly subscribed instance is returned, it will be moved to the TencentDB recycle bin and retained there for seven days. During the retention period, the instance cannot be accessed, but it can be restored after renewal. For specific operations, please refer to Renewal Instructions.
- After a pay-as-you-go instance is returned, it will be moved to the recycle bin and retained there for three days. During the retention period, the instance cannot be accessed, but it can be restored after a timely top-up.

Notes

After an instance is completely terminated, its data will be cleared and cannot be recovered. You need to back up the data in advance.

Notes

- When an instance is terminated, its read-only instances will not be terminated simultaneously.
- After an instance is terminated, it will be moved to the recycle bin. During the retention period, the instance cannot be accessed. If you want to restore the instance, you can do so in the recycle bin. For specific operation details, see Recycle Bin.
- When an instance is terminated, its IP resources will be released simultaneously, and its disaster recovery instance will stop the sync connection and be automatically promoted to primary instance.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB instance is in **Running** status.

Operation Steps

Monthly subscribed instance

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the **Operation** column of the target instance, select **More** > **Return and Refund**.

> **Note:**
> Note: For annual and monthly subscription instances, if the return button is unavailable, it indicates that the account has used up the self-service return limit for annual and monthly subscriptions. The instance cannot be manually destroyed and will be automatically destroyed upon expiration.

©2013-2025 Tencent Cloud. All rights reserved.

6. In the pop-up dialog box, click **View Details** to confirm the information of the instance to be destroyed.
7. Carefully read the precautions for destroying instances, check **I have read and agreed to**, confirm the destruction, and click **Yes**.

8. On the **Please Confirm the Following Refund Information** page, confirm the **Refund List** information, verify the refund amount, and click **Confirm Refund** if everything is correct.

Pay-As-You-Go Instance

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the **Operation** column of the target instance, select **More** > **Terminate**.
6. In the pop-up dialog box, read the prompt message carefully, confirm the instance to be terminated, and click **OK**.

Recycle Bin

Terminated instances will be put into the recycle bin and can be restored during the retention period as instructed in Recycle Bin.

©2013-2025 Tencent Cloud. All rights reserved.

Adjusting Oplog Capacity

Last updated: 2025-02-08 10:22:37

Overview

Oplog is a key component in MongoDB used to record database operation logs. The capacity of Oplog should be at least 10% of the node capacity. This is because Oplog records all database operations, including insertions, updates, and deletions. If the Oplog capacity is too small, it may be overwritten, affecting MongoDB's rollback feature. When purchasing an instance, the default Oplog size is 10% of the instance size, which can be expanded on-demand to 90% of the instance size. Reduction is not supported at this time.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- If it is an on-demand instance billing, please ensure your Tencent Cloud account balance is sufficient.
- The instance and its associated instances are in normal state (running) and are not currently executing any tasks.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the **oplog information** column, click **View/Adjust**.

6. In the **Adjusting Oplog** dialog box, confirm instance information and evaluate the capacity to be adjusted based on the current Oplog capacity.

> **Note:**
> Oplog capacity should be at least 10% of the node capacity. When the size of the oplog file reaches its maximum capacity, MongoDB will start overwriting previous operation records from the beginning of the file. If the Oplog is too small, data may be quickly overwritten, leading to operational risk.

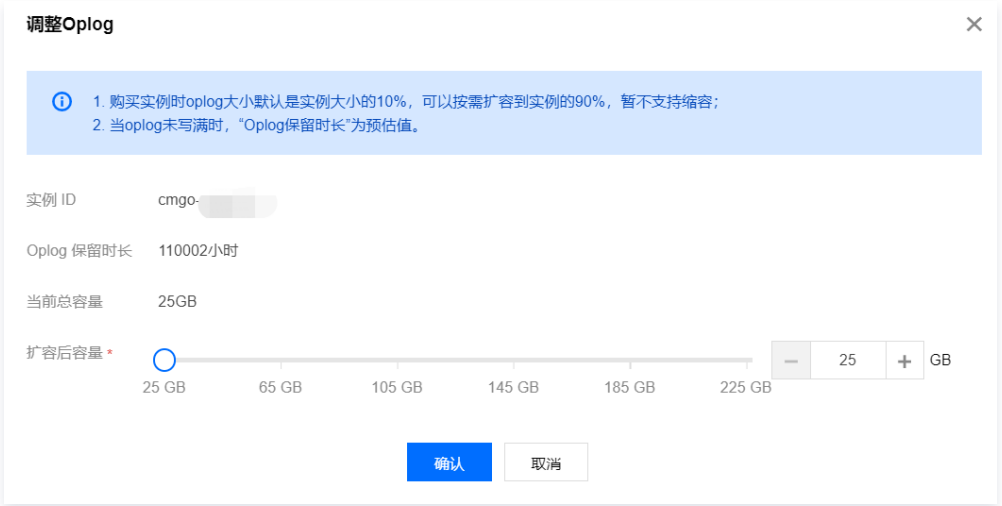

©2013-2025 Tencent Cloud. All rights reserved.

Interface Parameter	Parameter interpretation
Oplog retention duration	The time when the current Oplog storage capacity is filled. When the Oplog is not full, the Oplog retention duration is an estimated value.
Current total capacity	The capacity of each Mongod node storing the Oplog.
Capacity after expansion	Slide the slider to set the Oplog capacity after expansion.

7. Click **OK**, select **Task Management** in the left navigation bar, wait for task completion, and complete the expansion.

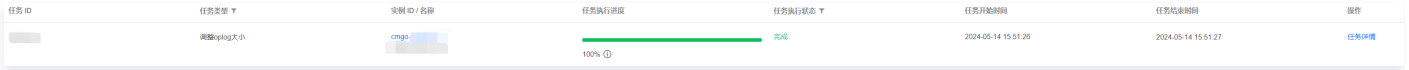

©2013–2025 Tencent Cloud. All rights reserved.

Node Management
Node Overview

Last updated: 2025-02-08 10:23:21

The replica set architecture of TencentDB for MongoDB achieves high availability and read/write separation by deploying multiple types of nodes. Each replica set instance consists of one primary node, one or multiple secondary nodes, and one hidden node. The sharded cluster architecture implements the horizontal capacity expansion of data based on the replica set architecture by combining multiple replica sets, each of which is a shard.
Each node is as described below:

Node	Function	Note:
primary node (Primary node)	It is responsible for executing and responding to data read/write requests.	There can be only one primary node in each replica set instance.
secondary node (Secondary node)	A secondary node replicates the data of the primary node by periodically polling its oplogs to ensure the data consistency. When the primary node fails, the system will elect a secondary node the new primary node in order to ensure the high availability.	• Client connections to the secondary node can only read data, not write data. • Secondary nodes support expansion. For details, see Adding Secondary Node. • Secondary nodes can be promoted to primary nodes. For detailed operations, see Promoting Replica Node to Primary Node.
Hidden node (Hidden node)	A secondary node will be designated as the hidden node by default for each newly purchased instance, serving as an invisible replica node to back up data. When a secondary node fails, it can be replaced with the hidden node as a new secondary node to achieve high availability.	• There can be only one hidden node in each replica set instance. • A secondary node set as the hidden node cannot be deleted. • The hidden node is not in the "primary node's standby list" and will not be elected as the primary node, but it will participate in the vote election for the primary node.
read-only node (ReadOnly node)	If the read-only replica feature is enabled, the system will set one or more secondary nodes as read-only nodes. They are mainly suitable for read request scenarios with an ultra high data volume. They sync data from the primary or secondary node through oplogs, and the system automatically routes read requests to them to reduce the access pressure on the primary node.	• Read-only nodes do not participate in the vote election for the primary node and will not be elected as the primary node. • A replica set instance can have multiple read-only nodes. For more information, see Add read-only nodes.

©2013-2025 Tencent Cloud. All rights reserved.

Viewing node information

Last updated: 2025-02-08 10:23:46

Overview

<TencentDB for MongoDB allows you to view the instance node information, including node ID, role, running status, used capacity, and other information. In addition, it supports node management operations, such as adjusting node specification, promoting replica node to primary node, enabling read-only replica, and configuring primary/replica failover. You can use node management to efficiently manage instance nodes and locate node-related exceptions during operation.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. View the mongod and mongos node information.

- **Mongod node**

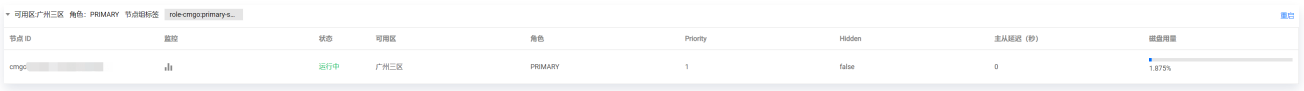

Parameter Name	Parameter interpretation
Node ID	<Mongod node ID number.>
Monitoring	Click , and you can view the monitoring views of various metrics of the node on the monitoring panel on the right. For more information, see Viewing Monitoring Data.
Status	The running status of the current node.
Availability Zone	The AZ of the region to which the current node belongs.
Role	The role of the current node. • PRIMARY: Primary node. • SECONDARY: Secondary node. • READONLY: Read-only node.
Priority	The priority of a node for being elected as the primary node. The greater the value, the higher the priority.
Hidden	Whether the node is hidden. Default value: `false`.
Master-slave latency (seconds)	The latency in syncing data from the primary node to the replica node in seconds.
Disk Utilization	Utilization of the node disk.

- **Mongos node**

Parameter Name	Parameter interpretation
Node ID	Mongos node ID number.
Monitoring	Click , and you can view the monitoring views of various metrics of the node on the monitoring

©2013-2025 Tencent Cloud. All rights reserved.

	panel on the right. For more information, see Viewing Monitoring Data.
Status	The operational status of the node.
Availability Zone	AZ of the mongos node.

©2013–2025 Tencent Cloud. All rights reserved.

Adjust the mongod node specification

Last updated: 2025-02-08 10:24:09

Overview

If the configuration of your purchased TencentDB for MongoDB instance isn't suitable for (either below or above) your current business requirements, you can quickly adjust its specifications based on your actual business conditions (at the initial stage, at the rapid development stage, during peak hours, or during off-peak hours), so as to better meet your needs such as full utilization of resources and real-time cost optimization.
Mongod configuration changes include: adjusting mongod's computing specifications and disk capacity. Before making changes, please familiarize yourself with the Product Specifications supported by TencentDB to help you choose specifications suitable for your business.

Billing Instructions

Adjusting instance configuration will start billing according to the new configuration. Please ensure your Tencent Cloud account balance is sufficient. For more information, see Billing Instructions for Configuration Change.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- If it is an on-demand instance billing, please ensure your Tencent Cloud account balance is sufficient.
- The instance and its associated instances are in normal state (running) and are not currently executing any tasks.

Expansion Principles and Impact

- When expanding storage space, if the physical machine resources of each node are sufficient, local expansion will be performed without cross-machine migration and switching, and there will be no connection interruption.
- When expanding storage space, if the physical machine resources of each node are insufficient, cross-machine data migration is required. After the expansion task initiation, the system will create nodes based on the required new specifications and synchronize data. Switching will be performed according to the switch time selected by the user, causing a connection fluctuation of about 10 seconds. It is recommended to implement a reconnection mechanism in your business code and adjust during business trough. When the data volume is large, the overall time consumption for configuration change will be longer.

> **Notes:**
> After the scaling specifications are executed, the system will trigger a new automatic backup task.

Notes

- During the adjustment process, there may be 1-2 instances of intermittent disconnection, each lasting about 10 seconds. It is recommended that the program has an automatic reconnection function.
- During the adjustment process, if you set the writeconcern concern level to write majority, there may be brief request delays. Please adjust the business timeout period accordingly.
- After the configuration adjustment, the instance name, intranet address, and port remain unchanged.
- Once the configuration adjustment task is initiated, you cannot cancel this operation midway.
- After the instance upgrade is completed, it is recommended to adjust the oplog capacity to avoid the oplog being too small, which may lead to data being overwritten and affect the rollback feature of MongoDB.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the **Operation** column of the target instance, select **Adjust Configuration** from the **Configuration Adjustment** dropdown list.

©2013-2025 Tencent Cloud. All rights reserved.

6. On the **Configuration Adjustment** page, you can readjust node memory and total node capacity. As shown below (using a shard as an example).

Parameter Name	Parameter interpretation	Parameter Example
Instance Name	Name of the instance with the configuration to be changed.	test-4dot2-XXX
Expiration Time	Expiration time of the instance, reminding the expiration time for instances billed on a yearly or monthly basis.	2022-04-24 19:23:43
Instance Architecture	Description of the cluster architecture of the instance. For details, see System Architecture.	Sharded cluster instance with 2 shards, each shard consists of a replica set of 3 storage nodes, totaling 6 storage nodes for the entire instance.
Current Node Memory/Total Capacity	The memory and total capacity of a single Mongod node in the current instance. For sharded clusters, the total capacity of a node is the capacity of a single shard's node. To query the number of CPU cores of an instance, refer to the Mongod specifications in the Product Specifications.	4GB/100GB
Node Memory	Re-select the memory of a single Mongod node in the dropdown list. For how to select specifications, refer to the Mongod specifications in the Product Specifications.	8GB
Total Node Capacity	Adjust the total disk capacity of a single Mongod node on the slider. For how to select specifications, refer to the Mongod specifications in the Product Specifications. **Notes:** When reducing disk capacity, it must be greater than or equal to 120% of the current used capacity.	500GB
Switch Time	• Select **Upon completion of adjustment** to immediately execute the switch instance specification task and switch to new specifications. Adjusting instance memory and disk may involve node migration or master-slave switch. The master-slave switch time point will be uncontrollable and may cause disconnection or restart.	Within the maintenance period

©2013-2025 Tencent Cloud. All rights reserved.

	• Select **Within the maintenance period** to execute the switch instance specification task within the maintenance period. For more information about maintenance time, see Set the instance maintenance time.	
Fees	• Pay-as-you-go: hourly billing price after instance configuration adjustment. Click **Billing Details** to view billing items, billing formula, and confirm the cost. • Annual and monthly subscription: for upgrade, show fees to be paid before the instance expires; for downgrade, show fees to be refunded before the instance expires. For billing details after configuration adjustment, see Billing instructions for configuration change.	1,201.83 CNY

7. After confirmation, click **Submit**.

Relevant API

API Name	Interface Feature Description
ModifyDBInstanceSpec	Modifying the configuration of TencentDB instance

©2013-2025 Tencent Cloud. All rights reserved.

Adding Secondary Node

Last updated: 2025-02-08 10:24:35

Overview

All replica nodes of the instance participate in system high availability support. In case of primary node failure, any secondary node may be elected as the new primary node to handle data write requests, thus the more replicas, the higher the availability. In read-heavy and write-light concurrent scenarios, enabling read-write separation can extend read performance through replica nodes, significantly reducing the read pressure on the primary node.
The total number of primary/secondary nodes in a TencentDB for MongoDB cluster supports 3, 5, or 7 nodes, meeting the configurations of 1 primary and 2 secondary nodes, 1 primary and 4 secondary nodes, and 1 primary and 6 secondary nodes, respectively. You can appropriately increase the number of secondary nodes based on actual business concurrency spikes. When the business volume decreases, reduce the number of secondary nodes to better achieve full utilization of resources and real-time cost optimization.

Billing Instructions

Adjusting instance configuration will start billing according to the new configuration. Please ensure your Tencent Cloud account balance is sufficient. For more information, see Billing Instructions for Configuration Change.

Notes

- New nodes join the cluster and start synchronizing data without affecting the business.
- Be sure to perform disaster recovery. It is recommended to initiate configuration change tasks during maintenance time. For more information on maintenance time, see Set the instance maintenance time.
- Do not initiate tasks to adjust the number of nodes, node computing specifications, and storage simultaneously.
- After adjusting the number of nodes, start billing according to the new specifications.
- After adjusting the number of nodes, the instance name, internal network address and port remain unchanged.
- Once the configuration adjustment task is initiated, you cannot cancel this operation midway.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- If it is an on-demand instance billing, please ensure your Tencent Cloud account balance is sufficient.
- The instance and its associated instances are in normal state (running) and are not currently executing any tasks.

Adding a Secondary Node (Replica Set)

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar's **MongoDB** drop-down list, select **Replica Set Instance**.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, in the **Mongod Node** tab, click **Adding Secondary Node**.
7. In the **Adjust Node Quantity** dialog box, review the considerations for adjusting the node quantity, confirm and configure the parameters according to the parameter explanation in the table below.

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Name	Parameter Explanation
Instance ID/Name	The instance name of the node quantity to be adjusted.
Expiration Time	Expiration time of the instance, reminding the expiration time for instances billed on a yearly or monthly basis.
Instance Configuration	Check the current specification of the instance, including the CPU core quantity, memory, disk capacity, and node quantity. The node quantity includes the total number of primary and secondary nodes. Evaluate the number of nodes to be added based on the current configuration.
Number of nodes to be added	Select the number of secondary nodes to be added in the drop-down list.
Deployment AZ	The AZ where all instance nodes are deployed. This parameter will be displayed if the instance nodes are in the same AZ.
secondary node-n	In the scenario where instances are in different AZs, this parameter indicates the different AZs corresponding to different instance nodes. The value range of n is 0 – 6. Configure the AZ for the newly added secondary node in the dropdown list.
Fees	• Pay-as-you-go: hourly billing price after instance configuration adjustment. Click **Billing Details** to view billing items, billing formula, and confirm the cost. • Monthly/Yearly Subscription: total cost within the expiration time after instance configuration adjustment. For billing details after configuration adjustment, see Configuration Adjustment Billing.

8. Confirm the fee information, click **OK** to complete the operation.
9. In the left navigation bar, select **Task Management** to view ongoing tasks, wait for the **task execution progress** to reach **100%**, and the **task execution status** to be **Completed**.

Adding a Single Shard Node (Shard Instance)

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Sharded Instance** from the **MongoDB** drop-down list.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, in the **Mongod Node** tab, click **Adding Secondary Node**.

©2013-2025 Tencent Cloud. All rights reserved.

7. In the **Adjust Single Shard Node Quantity** dialog box, review the considerations for adjusting the node quantity, and configure the parameters according to the parameter explanation in the table below.

- **The instance is deployed in the same AZ**

- **The instance is deployed in different AZs**

©2013–2025 Tencent Cloud. All rights reserved.

Interface Parameters	Parameter Explanation
Instance ID/Name	Please confirm the instance name for which the number of nodes per shard needs to be adjusted.
Instance Configuration	Please understand the current configuration specifications of the instance, including cpu cores, memory, disk capacity, and node count. The node count includes the total number of primary and secondary nodes. The number of nodes per shard is evenly distributed according to the shard count. For example, if there are 2 shards and 8 nodes, then each shard has 4 nodes. Please evaluate the number of additional nodes needed based on the current configuration.
Number of nodes to be added	In the dropdown list, select the number of secondary nodes to add for a single shard.

©2013–2025 Tencent Cloud. All rights reserved.

Deployment AZ	The AZ where all instance nodes are deployed. This parameter will be displayed if the instance nodes are in the same AZ.
secondary node-n	In scenarios where instances are in different AZs, this parameter indicates the different AZs of the instance's different shard nodes. The value range of n is 0 – 6. Please configure the corresponding AZ for the newly added secondary node in the dropdown list.
Fees	• Pay-as-you-go: The hourly billing price after the instance configuration is adjusted. Click billing details to view billing items, billing formula, and confirm cost information. • Annual and monthly subscription: The total cost within the expiration time after the instance configuration is adjusted. For billing details after configuration adjustment, see billing instructions for configuration change.

8. Confirm the fee information, click **OK** to complete the operation.
9. In the left navigation bar, select **Task Management** to view ongoing tasks, wait for the **task execution progress** to reach **100%,** and the **task execution status** to be **Completed**.

Relevant API

API Name	Interface Feature Description
ModifyDBInstanceSpec	Modifying the configuration of TencentDB instance

©2013-2025 Tencent Cloud. All rights reserved.

Deleting Secondary Node

Last updated: 2025-02-08 10:24:57

Overview

Deleting a secondary node will reduce the high availability of the cluster. When business volume is low, you can appropriately reduce the number of secondary nodes to avoid resource waste.

Must-Knows

- Deleting a secondary node will reduce the high availability of the cluster. Please operate with caution. Ensure that after deleting the secondary node, the total number of cluster nodes meets 3, 5, or 7, i.e., 3 nodes with 1 primary and 2 secondary, 5 nodes with 1 primary and 4 secondary, or 7 nodes with 1 primary and 6 secondary.
- Hidden nodes cannot be deleted. When a secondary node fails, the system will automatically switch it with the hidden node to ensure the high availability of the cluster.
- The IP address of the deleted secondary node will no longer be retained, causing the link to the secondary node to be disconnected.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- If it is an on-demand instance billing, please ensure your Tencent Cloud account balance is sufficient.
- The instance and its associated instances are in normal state (running) and are not currently executing any tasks.

Deleting a Secondary Node (Replica Set)

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar's **MongoDB** drop-down list, select **Replica Set Instance**.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, under the **Mongod Node** tab, in the instance list, select the secondary node information to be deleted, and choose **Node Operation > Delete Secondary Node**.

   > **Note:**
   > In the node management list, if the **Hidden** column of the node is **true**, it cannot be deleted. Note: Hidden nodes do not support deletion.

7. In the **Delete Secondary Node** dialog box, review the precautions for adjusting the number of nodes, and confirm the instance name, expiration time, etc. for the adjustment.

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Name	Parameter Explanation
Instance ID/Name	The instance name of the node quantity to be adjusted.
Expiration Time	Expiration time of the instance, reminding the expiration time for instances billed on a yearly or monthly basis.
Instance Configuration	Please understand the current configuration specifications of the instance, including cpu cores, memory, disk capacity, and node count. The node count includes the total number of primary and secondary nodes. The number of nodes per shard is evenly distributed according to the shard count. For example, if there are 2 shards and 8 nodes, then each shard has 4 nodes. Please evaluate the number of additional nodes needed based on the current configuration.
Delete Secondary Node Information	Please confirm the information of the node to be deleted, including: node ID, AZ, node role, and node Tag.
Configuration Change Fee	Fees after specification adjustment. In pay-as-you-go billing mode, fees are charged hourly by the new specification in three billing tiers. For annual and monthly subscription: the total cost of the remaining duration of the new specification.
Compare	You can compare the specification and maximum number of connections met before and after the change of Mongod secondary nodes to assess whether the new specification meets your needs.

8. Confirm cost information, click **OK**, and complete operation.
9. In the left navigation bar, select **Task Management** to view ongoing tasks. Wait for **task execution progress** to reach **100%,** and **execution status** to be **Completed**.

Relevant API

API Name	Interface Feature Description
ModifyDBInstanceSpec	Modifying the configuration of TencentDB instance

©2013-2025 Tencent Cloud. All rights reserved.

Add read-only nodes

Last updated: 2025-02-08 10:25:19

Overview

When your business has a massive number of read requests, it may be difficult for the primary and secondary database nodes to process such requests, causing a high latency, slow response, and severely dropped throughput of business requests. TencentDB for MongoDB provides read-only nodes with an independent connection address. They can sync data from a primary or secondary node with the lowest latency through oplog. You can create one or multiple read-only nodes to implement read/write separation and relieve the access pressure on the primary and secondary nodes.

Note:

- If a business only accesses read-only nodes, it is recommended to configure two or more read-only nodes to achieve load balancing for read requests with high availability guarantee. The connection string of the read-only node can be directly obtained from the network configuration on the **Instance Details** page.
- Read-only nodes are not in the "primary node's standby list", will not be elected as primary nodes, and will not participate in voting for primary node election.

Version Description

TencentDB for MongoDB 7.0, 6.0, 5.0, 4.4, 4.2, and 4.0 support adding read-only nodes, while 3.6 doesn't.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Mongod Node** tab, click **Add Read-only Node**.

Parameter Name	Parameter Explanation
Instance	Check the current specification of the instance, including the CPU core quantity, memory, disk capacity,

©2013-2025 Tencent Cloud. All rights reserved.

Configuration	total node quantity, and read-only node quantity to evaluate the number of read-only nodes to be added.
Add Read-Only Node	The number of new read-only nodes. Value range: 0-5.
deployment AZ	The AZ where all read-only nodes are deployed. This parameter will be displayed if the instance nodes are in the same AZ.
Compare	Compare the specifications before and after read-only nodes are added to assess whether the new specification meets your needs. • The specification of a replica set instance includes mongod specification, disk capacity, number of read-only nodes, and maximum number of connections. • The specification of a sharded cluster instance includes number of shards, mongod specification, disk capacity, number of read-only nodes, and maximum number of connections.
Total Fees	• Pay-as-you-go: The hourly billing price after the instance configuration is adjusted. Click billing details to view billing items, billing formula, and confirm cost information. • Annual and monthly subscription: The total cost within the expiration time after the instance configuration is adjusted. For billing details after configuration adjustment, see billing instructions for configuration change.

7. To confirm adding a read-only node, click **Yes**.
8. On the left sidebar, select **Task Management**. In the task list, find the instance by ID or name and wait for **Task Status** of the read-only node adding task to become **Completed**.

©2013-2025 Tencent Cloud. All rights reserved.

Adjust the shard quantity

Last updated: 2025-02-08 10:25:43

Overview

The number of shards in a sharded cluster can be adjusted according to business traffic after purchasing the instance to adapt to changing business scenarios.

Description

Adjusting instance configuration will start billing according to the new configuration. Please ensure your Tencent Cloud account balance is sufficient. For more information, see Billing Instructions for Configuration Change.

Notes

- New nodes join the cluster and start synchronizing data without affecting the business.
- Do not initiate tasks to adjust the number of nodes, node computing specifications, and storage simultaneously.
- After adjusting the number of nodes, start billing according to the new specifications.
- After adjusting the number of nodes, the instance name, internal network address and port remain unchanged.
- Once the configuration adjustment task is initiated, you cannot cancel this operation midway.
- Reducing the number of shards is not allowed.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- If it is an on-demand instance billing, please ensure your Tencent Cloud account balance is sufficient.
- The sharding instance and its associated instances are in normal status (running) and are not currently executing any tasks.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Sharded Instance** from the **MongoDB** drop-down list.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click its **Instance ID** to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, under the **Mongod Node** tab, click **Adjust Shard Count**.
7. In the **Adjust Shard Count** dialog box, review the precautions for adjusting shard count, refer to the table below, specify the new shard count, and confirm the cost.

©2013-2025 Tencent Cloud. All rights reserved.

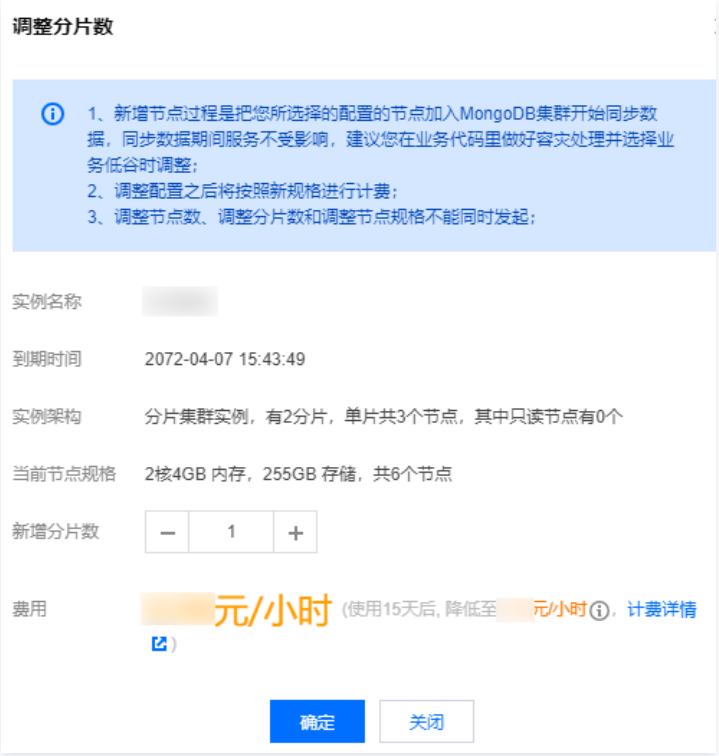

Parameter Name	Parameter interpretation	Parameter Example
Instance Name	The instance name of the node quantity to be adjusted.	test-4dot2-XXX
Expiration Time	Expiration time of the instance, reminding the expiration time for instances billed on a yearly or monthly basis.	2022-04-24 19:23:43
Instance Architecture	Description of the instance's cluster architecture. For details, see System Architecture.	Sharded cluster instance with 2 shards, each shard has 5 storage nodes.
Current Node Specifications	Specifications of a single shard node in the current sharded cluster instance, including cpu cores, memory, storage capacity, and number of nodes.	2-core 4GB memory, 250GB storage, 10 nodes in total.
Number of new shards	Select the number of additional shards for the instance, ranging from [current number of shards, 36].	3
Fees	• Pay-as-you-go: hourly billing price after instance configuration adjustment. Click **Billing Details** to view billing items, billing formula, and confirm the cost. • Annual and monthly subscription: increase the number of shards, display the required payment within the expiration time. For billing details after configuration adjustment, see Configuration Adjustment Billing.	6.69 CNY/hour

8. Confirm no errors, click **OK**.

Relevant API

API Name	Interface Feature Description
ModifyDBInstanceSpec	Modifying the configuration of TencentDB instance

©2013-2025 Tencent Cloud. All rights reserved.

Adjust the mongos node specification

Last updated: 2025-02-08 10:26:03

Overview

Upgrading the computing specification of mongos nodes can increase the maximum number of connections to the database. You can adjust the mongos node specification appropriately based on the actual conditions of your business access.

Must-Knows

Upgrading the CPU performance and memory capacity of mongos nodes may involve cross-node data migration and cause a momentary disconnection. Therefore, before performing this operation, make sure that your business has an automatic reconnection mechanism. We recommend that you complete this operation within the maintenance time during off-peak hours.

Version Description

MongoDB 4.0 and later versions support adjusting the specifications of mongos.

Prerequisites

- Instance type: Sharded cluster instance.
- Instance status: Running.
- The CPU performance and memory capacity of the mongos nodes are insufficient and need to be upgraded.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Sharded Instance** from the **MongoDB** drop-down list.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, click the **Mongos Node** tab.
7. On the Mongos Node page, click **Mongos Configuration Change**, and in the pop-up dialog box, configure the new mongos specifications.

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Name	Parameter Explanation
Instance ID/Name	The unique ID and name of the instance.
deployment AZ	The AZ of the instance.
Mongos Quantity	The current number of mongos nodes.
Mongos Specs	Select the new mongos node specification in the drop-down list, which can be 1-core 2 GB MEM, 2-core 4 GB MEM, 4-core 8 GB MEM, 8-core 16 GB MEM, or 16-core 32 GB MEM.
Switch Time	Select to execute the resize instance task immediately upon completion of the selection. Adjusting instance memory and capacity may involve node migration or master-slave switch, which can be uncontrollable and may cause disconnection or restart. If you select During maintenance time, the task will be executed during the maintenance time. For more information, see Setting Instance Maintenance Time.
Configuration Change Fee	Fees after specification adjustment. In pay-as-you-go billing mode, fees are charged hourly by the new specification in three billing tiers. For annual and monthly subscription: the total cost of the remaining duration of the new specification.
Compare	You can compare the maximum number of connections before and after the mongos specification adjustment to assess whether the new specification meets your needs.

8. To confirm change of this specification, click **OK**.

©2013-2025 Tencent Cloud. All rights reserved.

Adding Mongos Node

Last updated：2025-02-08 10:26:25

Overview

You can add more mongos nodes to increase the maximum number of connections to the database instance.

Version Description

The current MongoDB 4.0 and above versions support adjusting the specifications of mongos.

Must-Knows

Increasing the number of mongos nodes will automatically bind IP addresses to the new mongos nodes and enable the connection string to access mongos. You can directly copy the connection string in the network section of the **Instance Details** page.

Prerequisites

- Instance type: Sharded cluster instance.
- Instance status: Running.
- The CPU performance and memory capacity of the mongos nodes are insufficient and need to be upgraded.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Sharded Instance** from the **MongoDB** drop-down list.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, click the **Mongos Node** tab.
7. On the **Mongos Node** page, click **Adding Mongos Node**.

- **The instance nodes are in the same AZ**

©2013-2025 Tencent Cloud. All rights reserved.

- **The instance nodes are in different AZs**

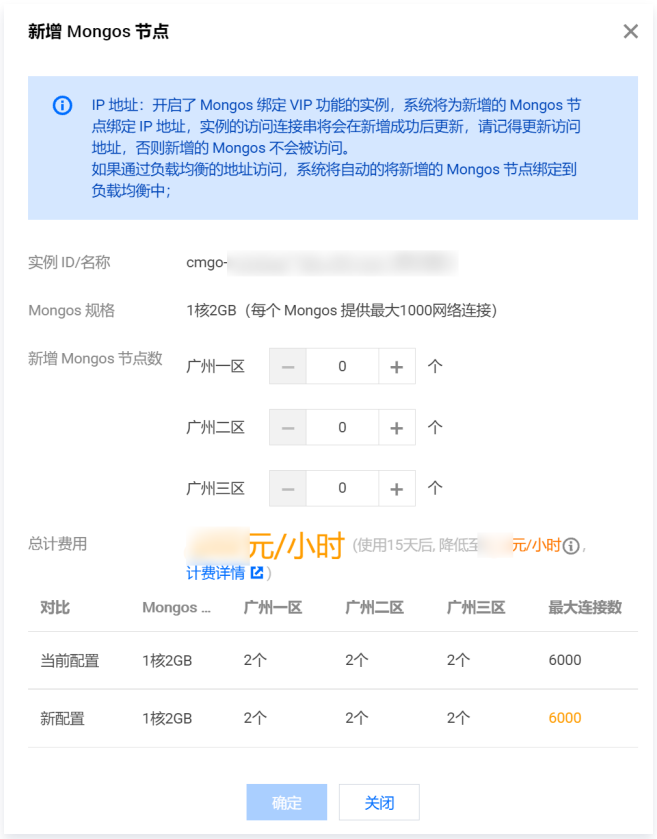

Parameter Name	Parameter Explanation
Instance ID/Name	The unique ID and name of the instance.
deployment AZ	This parameter indicates the AZ of the instance. It will be displayed if the instance is in the same AZ.
Mongos Quantity	The current number of mongos nodes configured for the instance. This parameter will be displayed if the instance nodes are in the same AZ.
Mongos Specs	Specification of mongos nodes, including the number of CPU cores, memory, and maximum number of connections.
Add Mongos Node	Select the number of mongos nodes to be added. An instance can have up to 48 mongos nodes.
Total Fees	Fees after configuration change. • In pay-as-you-go billing mode, fees are charged hourly by the new specification in three billing tiers. • Annual and monthly subscription: total fees for the remaining duration of the new specification.
Compare	You can compare the specification, number of nodes in the AZ, and maximum connection count before and after mongos nodes are added to assess whether the new specification meets your needs.

8. After confirmation, click **OK**.

©2013-2025 Tencent Cloud. All rights reserved.

Enabling Mongos Access Address

Last updated: 2025-02-08 10:26:45

Overview

After enabling the mongos access address of a sharded cluster instance, you can access the instance at this address. On the **Instance Details** page, you can see the mongos access connection string (for private network access).

Must-Knows

- Under the current VIP of the instance, different vports will be bound to different mongos nodes.
- After a mongos node fails, the system will bind its vport to a new mongos process, and the VIP and vport will remain unchanged.
- Enabling mongos access address won't affect the original CLB access address.

Version Description

MongoDB 4.0 and above support adjusting the specifications of mongos and enabling the mongos access address.

Prerequisites

- Instance type: Sharded cluster instance.
- Instance status: Running.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Sharded Instance** from the **MongoDB** drop-down list.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, click the **Mongos Node** tab.
7. On the **Mongos Node** tab, click **Enabling Mongos Access Address**.
8. In the pop-up dialog box, confirm the impact of enabling the access connection string and click **Yes**.

9. On the left navigation bar, select **Task Management**. In the task list, find the instance with the **Task Type** being **Enabling Node Access Address** by instance ID and wait for the **Task Execution Status** to become **Completed**.
10. On the left navigation bar, select **Sharded Instance**. In the instance list, find the instance with the access address enabled, click its instance ID to enter the **Instance Details** page. In the **Network Configuration** section under **Access Address**, you can view the Mongos access address. Hover over the connection string of the access address and click to directly copy the connection string for Mongos node access.

©2013-2025 Tencent Cloud. All rights reserved.

Mongos 访问地址：

连接类型	访问地址（连接串）
访问读写主节点	mongodb://mongouser:******@[redacted]rce=admin
仅读只读节点	mongodb://mongouser:******@[redacted]est? authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role-cmgo:readonly-group
仅读从节点	mongodb://mongouser:******@[redacted]t? authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role-cmgo:primary-secondary-group
仅读从节点和只读节点	mongodb://mongouser:******@[redacted]st? authSource=admin&readPreference=secondaryPreferred

©2013–2025 Tencent Cloud. All rights reserved.

Promoting Replica Node to Primary Node

Last updated: 2025-02-08 10:27:07

Overview

A TencentDB for MongoDB replica instance can have only one primary node but multiple replica nodes. When an anomaly is detected in the primary node, you can proactively promote a replica node to the primary node to ensure normal business operation. For sharded instances, all shard nodes and multiple replica nodes are divided into one primary node group and multiple secondary node groups. If some shard nodes in the primary node group are abnormal, you can proactively promote all nodes in the secondary node group to the primary node group.

Version Description

MongoDB version 3.2 and above supports promoting a replica node to the primary node.

Notes

Promoting to the primary node will disconnect the current TCP connections to the database. Please confirm that the business has an automatic reconnection mechanism before operation. Otherwise, you need to manually reconnect to the database using a new connection.

Prerequisites

Instance status: Running.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, in the **Mongod Node** tab, find the secondary node that needs to be promoted in the node list.
   - Replica Set: In the node list, find the secondary node that needs to be promoted, and click **Operation** in its **Promote to Primary Node** column.

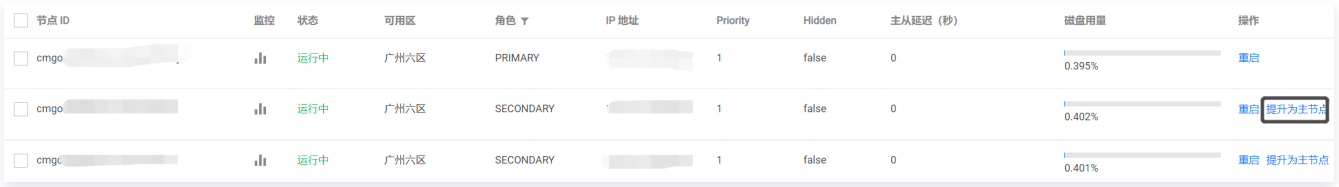

   - Sharded Cluster: In the node group list, find the secondary node group that needs to be promoted, and click **Promote to Primary Node** in the upper right corner.

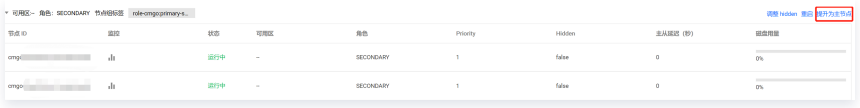

7. In the **Promote to Primary Node** dialog box, understand the impact of promoting the primary node, check **Confirm the risk of promoting to primary node**, and click **Yes**.

> **Note:**
> Promoting to primary node will cause the current TCP connections to the database to be disconnected. Please confirm that the business has an automatic reconnection mechanism before operation. Otherwise, you need to manually reconnect to the database.

©2013-2025 Tencent Cloud. All rights reserved.

8. Return to the **Instance Details** tab, check the **Instance State** for **Switching Primary Node**, and wait for the instance status to no longer have this information, indicating that the task is complete. On the **Node Management** page, you can see that the original secondary node's **Role** is PRIMARY.

Restarting a Node

©2013–2025 Tencent Cloud. All rights reserved.

Last updated: 2025-02-08 10:27:28

Overview

For daily operation and maintenance, restart the node to release memory or clear the cache and optimize resource usage; if the node fails, restarting may help recovery services.

Notes

- Restarting the node will disconnect existing connections. Please confirm that the application has a reconnection mechanism.
- Restarting the primary node will cause a master-slave switch and may result in data loss not yet synchronized to the secondary node. Please confirm the risk.

Prerequisites

Instance status: Running.

Restarting a Mongod Node

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left navigation bar, select Replica Set Instance or Sharded Instance.
3. On the right side, select a region at the top of the **Instance List** page.
4. In the Instance List, find the target instance whose node you want to view.
5. Click the instance ID to enter the **Instance Details** page, and click the **Node Management** tab.
6. On the **Node Management** page, go to the **Mongod Node** tab, and find the node to be restarted in the node list.
7. In the **Operation** column, click **Restart**.
   - Replica Set

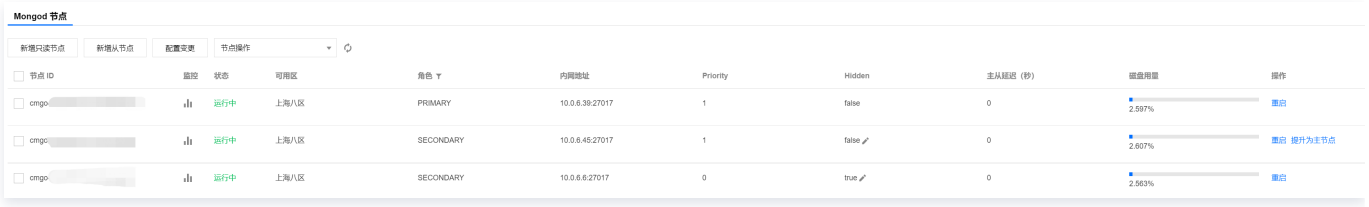

   - Sharded Instance

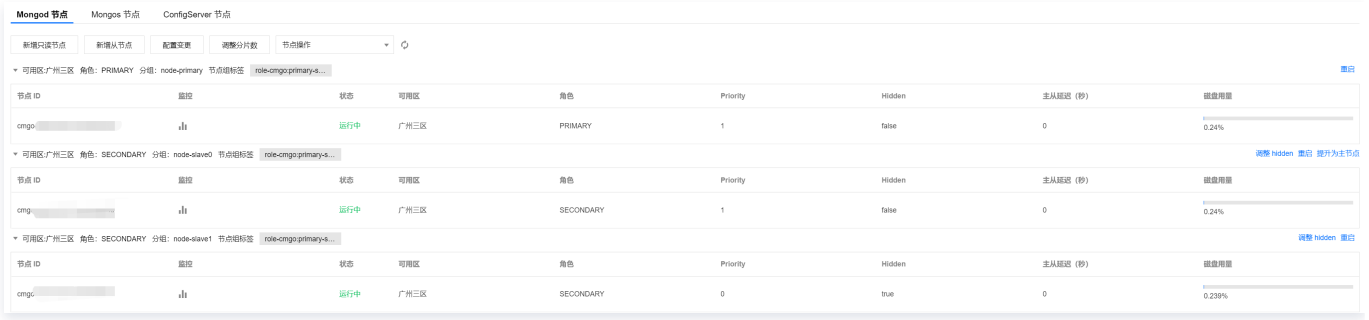

8. In the **Restarting a Mongod Node** pop-up window, review the precautions, check **Confirm the risk of restarting the node**, and click **Yes**.

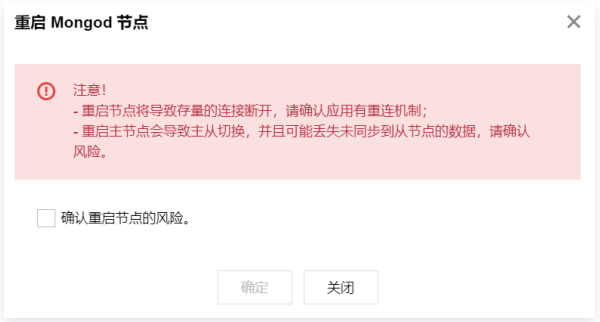

©2013-2025 Tencent Cloud. All rights reserved.

9. Wait for task completion.

Version upgrade

Last updated: 2025-02-08 10:27:52

©2013-2025 Tencent Cloud. All rights reserved.

Overview

TencentDB for MongoDB allows you to upgrade both the major and minor database versions. You can enjoy more features by upgrading the major version from v3.6 to v4.0 or from v4.0 to v4.2.

Version Description

- MongoDB supports upgrading from an old version to a higher version but does not support cross-version upgrades. It supports upgrading from MongoDB v3.6 to v4.0, v4.0 to v4.2, v4.2 to v4.4, and v4.4 to v5.0. For feature differences between versions, see Feature.
- You can also upgrade the minor version, for example, the minor version WT.40.3.34 of the major version 4.0.
- Minor version upgrade, the system will automatically detect the minor version and upgrade to the latest version, and you cannot select a target version.

> **Notes:**
> Cross-version upgrade, if you need to upgrade from version 3.2 to version 4.0, it can be achieved by using migration. For details, see Create Migration Task.

Notes

The upgrade process is completely automatic, and there will be several momentary interruptions during the process. We recommend that you upgrade during off-peak hours.

Prerequisites

- The instance is not a read-only or disaster recovery instance.
- Instances with the version to be upgraded are in normal state (running) and are not currently executing any tasks.
- The target version is confirmed.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the instance list, find the target instance.
5. In the **Instance ID/Name** column of the target instance, click the instance ID to enter the **Instance Details** page.
6. In the **Instance Details** page's **Specification Information** section, upgrade the major or minor version of the instance.
   - If the instance is version **3.6**, click **Upgrade to v4.0** after **Version and Engine** to upgrade the version from 3.6 to 4.0.
   - If the instance is version **4.0**, click **Upgrade to v4.2** after **Version and Engine** to upgrade the version from 4.0 to 4.2, and so on.
   - Click **Upgrade Minor Version** to upgrade the minor version to the latest version by default. For more information, see Versions and Storage Engines.

7. In the **Note** pop-up window, read the prompt message carefully, confirm the upgrade, and click **Yes**.
8. Wait for the upgrade task to complete. Once the instance status returns to **running**, you will see the version information has been updated.

©2013-2025 Tencent Cloud. All rights reserved.

Notes:
The upgrade is expected to complete within a few minutes. Please be patient.

Network Configuration

©2013–2025 Tencent Cloud. All rights reserved.

Switching Instance Network

Last updated: 2025-02-08 10:28:28

You can directly switch the network of a TencentDB for MongoDB instance in the console to adjust the network status promptly.

Background

Tencent Cloud supports classic network and Virtual Private Cloud (VPC), offering a diversity of high-quality services. On this basis, we provide more flexible services to help you manage network connectivity with ease.

- Switch from classic network to VPC: a single TencentDB source instance can be switched from classic network to VPC.
- Switch from VPC A to VPC B: A single TencentDB primary instance can be switched from VPC A to VPC B.

Version Description

Currently, MongoDB 3.2 and above support instance network switch.

Billing Instructions

Switching the database instance network doesn't generate additional costs.

Notes

- Switching the network may cause the change of instance's private IP. The original IP will become invalid after the repossession time has elapsed. Modify the instance IP on the client promptly.
- The switch from classic network to VPC is irreversible. After the switch to a VPC, the TencentDB instance cannot communicate with Tencent Cloud services in another VPC or classic network.
- After you switch the network of a primary instance, if there are read-only replicas or disaster recovery instances mounted, the networks of these read-only replicas or disaster recovery instances won't be automatically switched along with the primary instance. You need to manually switch their networks.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the target instance ID to enter the **Instance Details** page and click **Change Network** on the right of **Network**.
6. In the **Change Network** pop-up window, select a VPC and subnet in the current region in the drop-down list next to **Network**. If existing networks can't meet your requirements, you can click **Create VPC** or **Create Subnet** to create a network and select it.

©2013-2025 Tencent Cloud. All rights reserved.

- **New IP Assignment Method**: Select **Auto Assign** or **Designate address**.
  - **Auto Assign**: The system will automatically assign an available IP address based on the currently selected network environment.
  - **Designate address**: You can enter a specific IP address in the **New IPv4 Address** input box. The specified IP address must be unoccupied and within the IP range of the specified VPC.

    > **Note:**
    > - The destination VPC can only be selected from the VPC networks in the same region where MongoDB resides.
    > - We recommend that the VPC where the CVM instance resides should be selected; otherwise, the CVM instance will not be able to access MongoDB over the private network, unless a CCN instance is created between the two VPCs.

- **Old IP address**: Replica set instances support immediate release; for sharding instances, you can select the release time of the old IP address from the dropdown list, supporting immediate release, release after 1 day, release after 2 days, release after 3 days, and release after 7 days. The IP address will be released after the retention period.

  > **Note**
  > When selecting delayed release, there will be a transition period for the IP address change, referred to as "latency". During the latency period, the old IP address can still be connected, but the new IP address has already taken effect. Once the latency period ends, the old IP address will be reclaimed, and the related cleanup task will be initiated to remove the configurations and records associated with the old IP address. All network connections on the old address will be immediately disconnected, so please choose the release time carefully.

7. Confirm the network switch and click **OK**. Return to the instance details page, where you can view the new network of the instance.

©2013–2025 Tencent Cloud. All rights reserved.

Enabling Dedicated Public Network Access

Last updated: 2025-02-08 10:29:10

Tencent Cloud MongoDB supports internal and external network access. This article will introduce how to configure the external network access address in the console to achieve external network access to the MongoDB database, enabling more flexible and convenient database management.

Implementation Scheme

TencentDB for MongoDB enables external network access service through Cloud Load Balancer (CLB). Configure the CLB instance listening port in the TencentDB for MongoDB console. When the external network accesses the public IP address and port number of the CLB instance, the CLB instance will forward the request to the corresponding backend server. The backend server of the CLB will map the internal and external networks, automatically forwarding public network requests to the corresponding MongoDB internal server. For more information about CLB and its backend services, please refer to the Cloud Load Balancer Product Documentation.
As shown in the diagram below, public network users access the CLB through the IP address 192.168.17.6 and port number 80. The backend service of the CLB forwards the request to the MongoDB database internal IP address 10.0.0.1 and port number 27017 in the actual running environment. In this way, public network users can access the MongoDB database through the CLB.

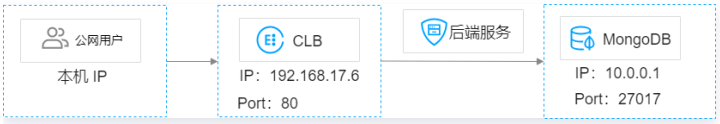

Use Limits

Before enabling the external network access feature of the MongoDB database, it is necessary to understand the related restrictions and requirements. These restrictions and requirements involve the MongoDB database, CLB (Cloud Load Balancer), network, etc., to ensure the security and stability of the database. For specific information, please refer to the table below.

Category	Category Subdivision	Restriction Description
TencentDB for MongoDB	Version	Supports MongoDB v4.0 and above.
	Sharded Cluster	Sharded clusters only support binding the instance's default access address (load balancing address) to CLB, and do not support binding individually opened Mongos addresses. **Note:** The load balancing address (LB address) of a MongoDB sharded cluster forwards client requests to the appropriate Mongos process for processing. For specific information, see System Architecture. The CLB listener will listen to the load balancer (LB) IP address and port number of the MongoDB sharded cluster.
	Replica Set	After adding or removing nodes in a MongoDB replica set, you need to modify the external IP address and specify the listening rules for the new nodes.
Network	Virtual Private Cloud	Only CLB instances under the same VPC as the cloud database MongoDB can be bound.
		After enabling external access, you cannot change the network of the database instance. To change it, you need to disable external access first.
	Security Group	MongoDB instances not bound to a security group do not support enabling external access. It is recommended to configure a security group to restrict incoming IP addresses. For specific operations, please see Configuring a Security Group.

©2013-2025 Tencent Cloud. All rights reserved.

Password Authentication	Password-Free Access	MongoDB instances with password-free access enabled do not support enabling external access.
CLB	Instance Type	MongoDB instances do not support binding to traditional CLB instances. For a comparison of the differences between Load Balancer (previously known as "Application Load Balancer") and Traditional Load Balancer, please see Instance Type Comparison.
	Instance Specifications	CLB instance specifications are divided into shared and performance capacity types. The maximum number of concurrent connections per minute for shared instances is 50,000. For some high-spec MongoDB instances, the performance of these shared instances may not meet the connection requirements. Therefore, it is recommended to choose an appropriate performance capacity CLB instance to meet your needs. For detailed differences between shared and performance capacity types, please see Instance Specifications Comparison.
	Account type	MongoDB instances do not support binding to CLB instances under traditional accounts, only under standard account types. For information on how to determine account types and account type upgrade methods, please refer to Account Type Description.
Limits of Operation	Close external network	Be sure to close external network access in the TencentDB for MongoDB console. Do not manually delete the listener created by MongoDB in CLB or delete the entire CLB instance, as this will cause business connection exceptions.
	Changing network	When changing the number of instance nodes, the external network feature may be affected. You need to update the external network configuration in the console to maintain smooth external network.

Notes

- After closing the external network access service, MongoDB will only remove the bound listener and will not release or recycle the CLB instance. The purchase and deletion of CLB instances are performed on the CLB side.
- It is recommended that each MongoDB instance dedicate a CLB instance. After binding, MongoDB will provide listener management and maintenance. If you need to share CLB instance with other resources, you must have a clear listener port management plan and reserve enough listeners. Otherwise, when multiple services use the load balancer, management confusion phenomenon may occur.

> **Note:**
> If the external network feature indicates a backend service health check exception, please go to the corresponding CLB console to check for health check risks or whether the health probe source IP segment is not allowed.

Prerequisites

- A TencentDB for MongoDB instance has been created, with version 4.0 or above, and is in normal operation.
- A CLB instance has been created and belongs to the same VPC as MongoDB, with normal operation.
- Please refer to MongoDB Compass Download (GUI) to download the visualization tools for Windows.

Directions

Step 1: Enabling External Network Service

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.

©2013-2025 Tencent Cloud. All rights reserved.

- If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.

5. In the **Instance ID/Name** column of the target instance, click **Instance ID** to enter the **Instance Details** page.
6. In the **Network Configuration** area of the **Instance Details** page, click **Public Network Access** and then **Configure CLB Public Network Access Service**.
7. In the **Service Authorization** dialog box, click **Grant**.

8. In the **Configure CLB Public Network Service** window, select the CLB listener instance and configure the listener rules.

a. In the **Bind CLB Instance** navigation tab, all CLB instances under the same VPC as the current MongoDB instance are listed. Please select the CLB instance to bind according to the required bandwidth limit specification. VIP refers to the public IP address of the CLB instance.

b. Click **Next**, in the **Configure Listener Rules** navigation tab, bind the CLB instance and set the monitoring rules.

- If it is a replica set, configure the CLB listening port number for the primary and secondary nodes of MongoDB.
- If it is a sharding instance, configure the corresponding listening port for the intranet address.

©2013–2025 Tencent Cloud. All rights reserved.

9. Click **OK**, and wait for the task to complete. In the network configuration area of the instance details page, you can view the connection string of the external access address.


网络配置

所属网络：bro-cd2 更换网络

所在子网：AutoName_20230704_100246

外网访问 ⓘ：已开启 修改 关闭

访问地址：

连接类型	访问地址（连接串）	外网访问地址（连接串）
访问读写主节点	mongodb://mongouser:******@ test?replicaSet=cmgo- &authSource=admin	mongodb://mongouser:******@ est?replicaSet=cmgo uthSource=admin
仅读从节点	mongodb://mongouser:******@ test?replicaSet=cmgo- &authSource=admin&readPreference=secondaryPreferred	mongodb://mongouser:******@ /test?replicaSet=cmgo- &authSource=admin&readPreference=secondaryPreferred

Log in to the CLB Console, in the instance list of instance management, find the CLB instance bound to MongoDB, click its **Instance ID**, enter the **Basic Information** tab of the instance, and select the listener management tab to view the corresponding listener.

Step 2: Configuring a Security Group

After enabling external network service, promptly configure security group rules for the CLB and its MongoDB instance to control access sources and ensure data access security.

1. Log in to the security group page of the Cloud Virtual Machine Console, create two security groups, the CLB instance security group and the MongoDB instance security group, and set **inbound rules** for each. For detailed steps, see Creating Security Group.

> **Note:**
> Security group **inbound rules** requirements are as follows:
> - CLB security group requirements: Allow client external IP address and listener port.
> - MongoDB security group requirements: Allow client external IP address and port 27017; allow the VIP address of the CLB instance, protocol port ALL (for health check).

2. Log in to the CLB Console, find the CLB instance bound to MongoDB in the instance list under instance management, click its **Instance ID**, go to the **Basic Information** tab, and select the **Security Group tab**.

   In the **Bound Security Group** area, click **Bind**, in the pop-up **Configure Security Group** window, select the security group of the created CLB instance, and click **Yes**. For detailed steps, see Configuring CLB Security Group.

©2013–2025 Tencent Cloud. All rights reserved.

3. Log in to the TencentDB for MongoDB console, find the instance to bind the security group in the instance list. Click the target **Instance ID**, select the **Data Security** tab, and click **Configure Security Group**. In the **Configure Security Group** dialog box, select the created MongoDB security group. Click **Yes**. For detailed steps, see Configuring a Security Group.

Step 3: Connecting to the Database Instance

1. On the TencentDB for MongoDB console **Instance Details** page, in the **Network Configuration** area, copy the connection string from the **Access Address** column under **Public Network Address (Connection String)** for either **Primary Node** or **Secondary Node**.

网络配置

所属网络：bro-cd2 更换网络

所在子网：AutoName_20230704_100246

外网访问 ⓘ：已开启 修改 关闭

访问地址：

连接类型	访问地址（连接串）	外网访问地址（连接串）
访问读写主节点	mongodb://mongouser:******@ test?replicaSet=cmgo- &authSource=admin	mongodb://mongouser:******@ est?replicaSet=cmgo uthSource=admin
仅读从节点	mongodb://mongouser:******@ test?replicaSet=cmgo- &authSource=admin&readPreference=secondaryPrefe rred	mongodb://mongouser:******@ /test?replicaSet=cmgo- &authSource=admin&readPreference=secondaryPrefe rred

2. Log in to the MongoDB Compass Download (GUI) client, paste the copied public network address connection string into the **URI** input box. The password information in the connection string is hidden as ***, **you need to manually replace** * with the instance's access password, and click **Connect**. As shown below:

©2013–2025 Tencent Cloud. All rights reserved.

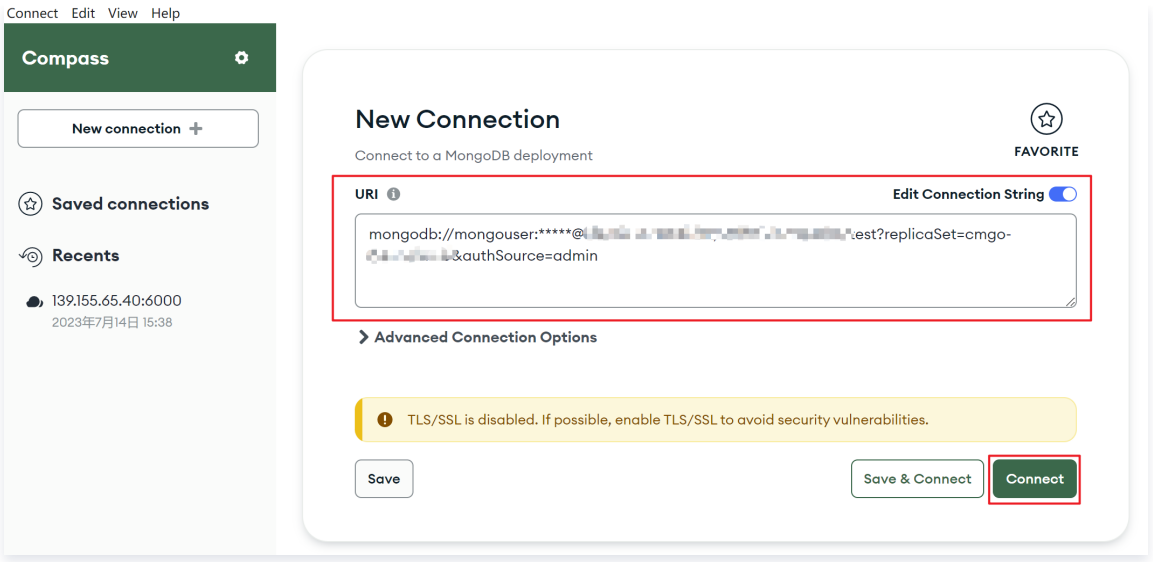

3. Wait for the connection to succeed, then you can manage the database in the client, as shown below:

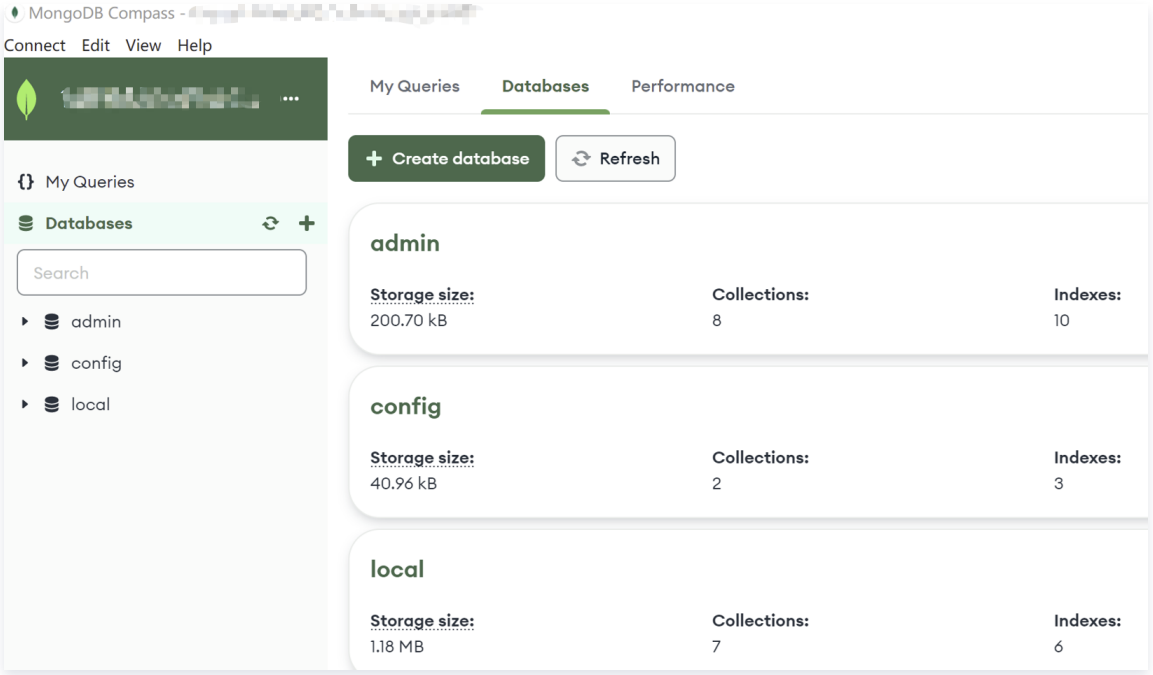

©2013–2025 Tencent Cloud. All rights reserved.

Enabling Quick Public Network Access

Last updated: 2025-02-08 10:30:49

The quick public network access feature of TencentDB for MongoDB provides users with a fast and convenient way to access the database externally. By enabling the quick public network feature with one click, the system will automatically assign an external network address and port for the instance, allowing users to easily access the database instance directly through the external network address without complex configuration.

> **Notes:**
> Note: The express external network is free, only for development, operation, testing, and debugging use, and does not guarantee data transmission security and bandwidth.

Version

MongoDB 4.0, 4.2, 4.4 & 5.0

Use Limits

- Sharded cluster only provides external network access address of lb address.
- The replica set provides external access addresses for primary, secondary, and read-only nodes (does not provide addresses for hidden nodes).
- Only one of dedicated and quick external network access can be enabled at a time; if you need to switch, you need to close it first and then reopen.
- A single MongoDB instance quick external network supports up to 1MB/s bandwidth, single CLB instance external network bandwidth limit: 20MB/s.
- A single MongoDB instance quick external network supports up to 200 concurrent connections.
- Quick external network access does not support security groups, only independent allowlists.
- Quick external network feature is not supported for password-free instances.

Notes

Replica set can support automatic primary-secondary switch, adding and removing nodes, switching network, switching AZ, and other node operations (automatically ensuring listener binding/unbinding).
• Master-slave switch external network binding follows VIP drift;
• Automatically apply for a new external network address when adding a node;
• Delete the bound address when deleting a node;
• Keep the external network address binding when switching networks and AZs;
• Automatically recycle the listener when destroying the instance;

Prerequisites

- A TencentDB for MongoDB instance has been created, with version 4.0 or above, and is in normal operation.
- A CLB instance has been created and belongs to the same VPC as MongoDB, with normal operation.

Enabling Quick Public Network Access

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.

©2013-2025 Tencent Cloud. All rights reserved.

5. In the **Instance ID/Name** column of the target instance, click **Instance ID** to enter the **Instance Details** page.
6. In the **Network Configuration** section, click **Public Network Access** and then **Enable Convenient Public Network Access**.

7. In the **Enable Quick External Network Access Service** window, edit the allowlist of IP addresses for external network access in the input box of **Access Permission**. The format requirements for IP addresses or IP address segments are as follows:
   - Single IP address, such as 10.23.12.24;
   - CIDR format: such as 10.23.12.24/24, where /24 indicates the length of the prefix in the address, range [24, 32];
   - Multiple IP addresses or IP address segments separated by English comma (,);
   - Allowing all addresses is not supported: 0.0.0.0/0

> **Notes:**
> After enabling, you can modify the allowlist address. In the **Network Configuration** area, click **Modify** to reset the allowlist address.

8. Click **Yes** and wait for the task to complete to access the external network.

Disabling Quick Public Network Access

After enabling quick public network access, if you need to disable this feature, you can go to the **Instance Details** page in the console, find the **Network Configuration** area, click **Off**, in the **Disabling Quick Public Network Access** window, click **Yes** to disable the feature.

©2013–2025 Tencent Cloud. All rights reserved.

FAQs

Last updated: 2025-02-08 10:31:12

The external network access configuration is managed by MongoDB, but users may independently operate the Cloud Load Balancer (CLB), causing management confusion and external network access interruption. To prevent this situation from occurring, the console will report some common issues.

Issue One: Cloud Load Balancer (CLB) Instance Mistakenly Deleted, Causing External Network Connection Failure

Phenomenon Description

MongoDB console prompts: "Cloud Load Balancer instance does not exist, please go to the Cloud Load Balancer console to check instance status."

Possible Reasons

After enabling the external network, the CLB instance bound by MongoDB suddenly does not exist and is mistakenly deleted.

Solution

1. Log in to the CLB Console. The previously bound instance no longer exists. Decide whether to recreate it according to needs.
2. Log in to the MongoDB Console, in the **Instance Details** section under **Network Configuration**, click **Public Network Access** and then **Off** to close the external network service.
3. Wait for the close task to complete, then click **Configure CLB Public Network Access Service**, select the CLB instance, configure the listening port, and reopen the external network service.

Issue 2: Listener Does Not Exist, Causing External Network Connection Not Working

Phenomenon Description

The MongoDB console prompt: "Listener does not exist, please go to the CLB Console to check the listener status."

Possible Reasons

After enabling external access, the user may have deleted the listener configuration on the CLB side.

Solution

1. Log in to the CLB Console, click **Instance ID**, go to the **Instance Details** page, then select the **Listener Management** page to confirm that the listener has been accidentally deleted.

2. Log in to the MongoDB Console, in the **Instance Details** section under **Network Configuration**, click **Public Network Access** and then **Off** to close the external network service.
3. Wait for the task to complete, then click **Configure CLB External Access Service**, select the CLB instance, reconfigure the listening port, and enable the external network service.

©2013-2025 Tencent Cloud. All rights reserved.

Issue 3: The Number of CLB Listeners Does Not Match the Number of MongoDB Node VIPs, Causing External Network Connection Failure

Phenomenon Description

MongoDB console prompt: "The number of listeners does not match the number of instance VIPs. If the current instance has added or removed nodes or deleted listeners, please click the modify button to configure the corresponding external access rules."

Possible Reasons

The listening port on the CLB side corresponds to the internal network address of each node of the MongoDB instance, as shown below. When the MongoDB instance adds a secondary node, the new node cannot find the corresponding listener, causing access error.

Solution

Log in to the MongoDB Console, in the **Instance Details** section of the **Network Configuration** area, click **Public Network Access** next to **Modify**, in the **Edit CLB External Service** window, modify the public network configuration. As shown below.

©2013–2025 Tencent Cloud. All rights reserved.

Issue 4: The Actual Listening Port of CLB Differs from the Configured Port, Please Check the Listening Rules of the CLB Instance or Modify the External Network Access Rules

Phenomenon Description

The MongoDB console prompts: "The actual listening port of CLB differs from the configured port, please check the listening rules of the CLB instance or modify the external network access rules."

Possible Reasons

The user mistakenly modified the IP binding port of the listener on the CLB side, causing the listening port to be inconsistent with the actual configured port.

Solution

Method 1: Log in to the MongoDB Console, in the **Instance Details** section of the **Network Configuration** area, click **Public Network Access** next to **Modify**, in the **Edit CLB External Service** window, modify the public network listening configuration port to match the listening port of the listener.
Method 2: Log in to the MongoDB Console, in the **Instance Details** section of the **Network Configuration** area, click **Public Network Access** next to **Off** to close the external network service. Wait for the closing task to complete, then click **Configure CLB Public Network Access Service**, select the CLB instance, reconfigure the listening port, and enable the external network service.

©2013–2025 Tencent Cloud. All rights reserved.

Space Analysis

Last updated：2025-02-08 10:31:40

Overview

TencentDB for MongoDB supports connection to DBbrain to analyze the computing and storage specifications of the mongod node. You can view the instance space utilization, including the sizes of data and logs, the daily increase in space utilization, the estimated number of available days, and the space used by the tables and databases of the instance.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharded instance is in **Running** status.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the instance list, find the instance for space analysis.
5. In the **Instance ID/Name** column of the target instance, click the instance ID in blue font to enter the **Instance Details** page.
6. In the **Instance Details** page's **Configuration Info** section, click **Space Analysis** after **mongod Specs**.

7. On the space analysis page in the DBbrain console, analyze the space usage details of the mongod node. For more information, see Space Analysis.

©2013-2025 Tencent Cloud. All rights reserved.

System monitoring
Monitoring Overview

Last updated: 2025-02-08 10:32:30

The monitoring feature provided by TencentDB for MongoDB allows real-time viewing of instance resource monitoring metrics. It supports visual graphs, tables, large screens, and various ways to count monitoring data. You can set alarm specifications and receive message notifications to help you understand database service anomalies in the first time, make timely business adjustments, and ensure stable business operations.

Monitoring granularity

TencentDB for MongoDB currently does not support custom selection of monitoring data collection granularity. The adaptive strategy is shown in the table below:

Time Span	Monitoring granularity	Retention Duration
0 to 1 day	5 second	1 day
0 to 1 day	1 minute	15 days
0 to 1 day	5 minutes	31 days
0 to 1 day	1 hour	93 days
0 to 1 day	1 day	186 days
0-7 days	1 hour	93 days
0-7 days	1 day	186 days
7-30 days	1 hour	93 days
7-30 days	1 day	186 days

Monitoring Description

TencentDB for MongoDB's replica set architecture achieves high availability capability by deploying multiple servers to store data replicas. Each replica set instance consists of a Primary node and one or more Secondary nodes. TencentDB for MongoDB's sharded cluster architecture implements horizontal data scaling based on the replica set architecture by combining multiple replica sets. Each sharded cluster instance is composed of mongos nodes, config server nodes, shard nodes, and other components. For more information, see System Architecture.
TencentDB for MongoDB monitors multiple performance metrics based on the cluster architecture and the request dispatch and return path, including resource utilization, request latency, request volume statistics, and network traffic. For more information, see the table below.

Monitoring Type	Monitoring Scope		
	Instance (master instance, read-only instance, disaster recovery instance)	Node	
		Mongod	Mongos
Resource Monitoring	Statistics of maximum CPU and memory usage, average usage rate, and disk usage of all nodes in the entire cluster.	Statistics of CPU, memory, and disk usage of Mongod nodes storing documents in the cluster.	Statistics of CPU and memory usage of Mongos nodes storing documents in the cluster.
Network Monitoring	Statistics of inbound and outbound traffic bytes of the entire cluster and the number of client connections.	Statistics of inbound and outbound bytes of Mongod nodes.	Statistics of inbound and outbound bytes of Mongos nodes.

©2013-2025 Tencent Cloud. All rights reserved.

Latency monitoring	Statistics of the time taken from the issuance of a request to the cluster to the final return.	Statistics of the average time taken for a request to reach a Mongod node and return from Mongod.	Statistics of the average time taken for a request to reach a Mongos node, be routed to Mongod by Mongos, processed by Mongod, and finally return from Mongos.
Request Monitoring	Statistics of the total number of requests issued to the cluster and the number of requests per second.	Statistics of the number of requests to access Mongod and the number of requests per second.	Statistics of the number of requests received by a Mongos node and the number of requests received per second.
Kernel monitoring	–	Statistics include: number of active read/write requests, read/write queue length, TTL, master–slave latency, cache hit rate, etc.	–

Monitoring Metrics

For all performance monitoring metrics and their meanings of TencentDB for MongoDB, please refer to the table below.

Instance

Monitoring dimension	Metric Name in Chinese	Monitoring Metric Name in English	Unit	Metric Description
CPU monitoring	Max Mongod CPU utilization	mongod_max_cpu_usage	%	The highest CPU utilization among all Mongod nodes in the cluster
	Average Mongod CPU utilization	monogd_avg_cpu_usage	%	Average CPU utilization of all Mongod nodes in the cluster
	Max Mongos CPU utilization	monogs_max_cpu_usage	%	The highest CPU utilization among all Mongos nodes in the sharded cluster
	Average Mongos CPU utilization	monogs_avg_cpu_usage	%	Average CPU utilization of all Mongos nodes in the sharded cluster
Memory monitoring	Max Mongod memory utilization	mongod_max_mem_usage	%	The highest memory utilization among all Mongod nodes in the cluster
	Average Mongod memory utilization	mongod_avg_mem_usage	%	Average memory utilization of all Mongod nodes in the cluster
	Max Mongos memory utilization	mongos_max_mem_usage	%	The highest memory utilization among all Mongos nodes in the sharded cluster
	Average Mongos memory utilization	mongos_avg_mem_usage	%	Average memory utilization of all Mongos nodes in the sharded cluster
Disk monitoring	Disk space utilization	disk_usage	%	Ratio of actual disk usage to the allocated disk space

©2013–2025 Tencent Cloud. All rights reserved.

Network Monitoring	Connection Count	cluster_conn	Count	Number of TCP connections to the instance.
	Connection percentage	connper	%	Proportion of current cluster connections to the maximum number of connections.
	Inbound traffic	cluster_view	Bytes	Statistics of inbound traffic bytes of the cluster.
	Outbound traffic	cluster_netout	Bytes	Statistics of outbound traffic bytes of the cluster.
Latency monitoring	Average latency of all requests	avg_all_request_delay	ms	Average latency of all requests executed in the cluster.
	Update average latency	avg_update_delay	ms	Average latency of cluster update requests
	Insert average latency	avg_insert_delay	ms	Average latency of cluster insert requests
	Average read latency	avg_read_delay	ms	Average latency of cluster read requests
	Average latency of aggregate requests	avg_aggregate_delay	ms	Average latency of cluster aggregate requests
	Average latency of Count	avg_count_delay	ms	Average latency of cluster Count requests. Count is used to count the number of documents in a collection that meet specified conditions.
	Average latency of Getmore	avg_getmore_delay	ms	Average latency of cluster Getmore requests
	Average latency of delete	avg_delete_delay	ms	Average latency of cluster delete requests
	Average latency of Command	avg_command_delay	ms	The average latency of command requests in the cluster. Command refers to all commands except insert, update, delete, and query.
	10ms – 50ms	10ms	Times	The number of requests executed between 10 milliseconds and 50 milliseconds.
	50ms – 100ms	50ms	Times	The number of requests executed in 50ms to 100ms.
	100ms	100ms	Times	The number of requests executed in over 100ms.
Request Monitoring	Total Requests	success_per_second	Times/second	The number of all requests successfully executed per second by the cluster.
	Insert Requests	insert_per_second	Times/second	The number of insert requests executed per second by the cluster. Insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.

©2013–2025 Tencent Cloud. All rights reserved.

	Read Requests	read_per_second	Times/second	The number of read requests executed per second by the cluster.
	Update requests	update_per_second	Times/second	The number of update requests executed per second by the cluster.
	Delete requests	delete_per_second	Times/second	The number of delete requests executed per second by the cluster.
	count request	count_per_second	Times/second	The number of count requests received per second by the cluster.
	Getmore request	getmore_per_second	Times/second	The number of Getmore requests received per second by the cluster.
	Aggregates requests	aggregate_per_second	Times/second	The number of aggregate requests per second by the cluster.
	Command requests	command_per_second	Times/second	The number of command requests received per second by the cluster. Command refers to all commands except insert, update, delete, and query.
Requests	Total Requests	node_success	Times	The number of all requests by the cluster.
	Insert Requests	node_inserts	Times	The number of insert requests received by the cluster. Insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.
	Read Requests	node_reads	Times	The number of read requests received by the cluster.
	Update requests	node_updates	Times	The number of update requests in the cluster.
	Delete requests	node_deletes	Times	The number of delete requests in the cluster.
	count request	node_counts	Times	The number of count requests received by the cluster.
	Getmore request	node_getmores	Times	The number of Getmore requests received by the cluster.
	Aggregates requests	node_aggregates	Times	The number of all requests aggregated by the cluster.
	Command requests	node_commands	Times	The number of command requests received by the cluster. Command refers to all commands except insert, update, delete, and query.

Mongod node

©2013–2025 Tencent Cloud. All rights reserved.

Monitoring dimension	Metric Name in Chinese	Monitoring Metric Name in English	Unit	Metric Description
CPU monitoring	CPU Uitilization	cpuusage	%	Percentage of CPU time used by processes on the Mongod node.
Memory monitoring	Memory Usage	memusage	%	Percentage of used memory space out of total memory capacity on the Mongod node.
Disk monitoring	Hard Disk Space Usage	diskusage	MBytes	Percentage of used disk space out of total space on the Mongod node.
	Number of disk reads	ioread	Times/second	Number of disk reads per second on a Mongod node.
	Number of disk writes	iowrite	Times/second	Number of disk reads per second on a Mongod node.
Network Monitoring	Inbound traffic	netin	Bytes	Statistics of inbound traffic bytes on a Mongod node.
	Outbound traffic	netout	Bytes	Statistics of outbound traffic bytes on a Mongod node.
Average request latency monitoring	Average latency of all requests	node_avg_all_requests_delay	ms	Average latency of all requests received by Mongod nodes.
	Update average latency	node_avg_update_delay	ms	Average latency of update requests on Mongod nodes.
	Insert average latency	node_avg_insert_delay	ms	Average latency of insert requests on Mongod nodes.
	Average read latency	node_avg_read_delay	ms	Average latency of read requests on Mongod nodes.
	Average latency of aggregate requests	node_avg_aggregate_delay	ms	Average latency of aggregate requests on Mongod nodes.
	Average latency of Count	node_avg_count_delay	ms	Average latency of count requests on Mongod nodes.
	Average latency of Getmore	node_avg_getmore_delay	ms	Average latency of getmore requests on Mongod nodes.
	Average latency of delete	node_avg_delete_delay	ms	Average latency of delete requests on Mongod nodes.
	Average latency of Command	node_avg_command_delay	ms	Average latency of command requests on Mongod nodes.

©2013–2025 Tencent Cloud. All rights reserved.

	10–50 milliseconds	10ms	Times	The number of requests executed between 10 milliseconds and 50 milliseconds.
	50–100 milliseconds	50ms	Times	The number of requests executed in 50ms to 100ms.
	100ms	100ms	Times	The number of requests executed in over 100ms.
Request Monitoring	Total Requests	node_success_per_second	Times/second	The number of all requests per second on a Mongod node.
	Insert Requests	node_insert_per_second	Times/second	The number of insert requests per second on a Mongod node. Insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.
	Read Requests	node_read_per_second	Times/second	The number of read requests per second on a Mongod node.
	Update requests	node_update_per_second	Times/second	The number of update requests per second on a Mongod node.
	Delete requests	node_delete_per_second	Times/second	The number of delete requests per second on a Mongod node.
	Count Request	node_count_per_second	Times/second	The number of count requests received per second on a Mongod node.
	Getmore request	node_getmore_per_second	Times/second	The number of getmore requests received per second on a Mongod node.
	Aggregates requests	node_aggregate_per_second	Times/second	The number of aggregate requests per second on a Mongod node.
	Command requests	node_command_per_second	Times/second	The number of command requests received per second on a Mongod node. Command refers to all commands except insert, update, delete, and query.
Kernel monitoring	Active write requests	aw	Count	Number of write requests in memory on a Mongod node.
	Active read requests	ar	Count	Number of read requests in memory on a Mongod node.
	Queued read requests	qr	Count	number of waiting read requests in the queue.
	Queued write requests	qw	Count	number of waiting write requests in the queue.

©2013–2025 Tencent Cloud. All rights reserved.

	Number of TTL deleted data entries	ttl_deleted	Times	Number of data entries automatically deleted by the database after TTL expiration.
	Number of TTL initiations	ttl_pass	Times	Number of data checks within the TTL time set in the database.
	Number of active sessions	active_session	Count	A session represents a single interaction between a client and a server. After establishing a connection, a session can be created for data read and write operations. Once created, the session remains active until the client closes the connection or it times out. This metric monitors the number of active sessions on the current Mongod node.
	Oplog retention duration	node_oplog_reserved_time	Hour	Oplog is used to record the operation logs of the database, and this metric tracks its retention duration.
	Source-Replica Delay	node_slavedelay	Second	A secondary node periodically polls the primary node's oplog to replicate the data of the primary node, and this metric tracks the latency of master-slave data synchronization.
	Cache hit ratio	replicaset_node	%	Current cluster cache hit ratio.
	Cache usage percentage	node_cache_used	%	Percentage of cache usage to total.
	Percentage of dirty cache data.	node_cache_dirty	%	Percentage of dirty cache data to total.
Requests	Total Requests	node_success	Times	Total number of requests in the cluster.
	Number of insert requests	node_inserts	Times	The number of insert requests in the cluster. Insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.
	Number of read requests	node_reads	Times	The number of read requests in the cluster.
	Number of update requests	replicaset_node	Times	The number of update requests in the cluster.
	Number of delete requests	node_deletes	Times	The number of delete requests in the cluster.
	Number of count requests	node_counts	Times	The number of count requests received by the cluster.
	Getmore requests	node_getmores	Times	The number of Getmore requests received by the cluster.
	Aggregates requests	node_aggregates	Times	The number of aggregate requests in the cluster.
	Command requests	node_commands	Times	The number of command requests received by the cluster. Command refers to all commands except insert, update, delete, and query.

Mongos Node (Sharded Cluster)

©2013-2025 Tencent Cloud. All rights reserved.

Monitoring dimension	Metric Name in Chinese	Metric Name in Chinese	Unit	Metric Description
CPU monitoring	CPU Uitilization	cpuusage	%	CPU utilization of Mongos nodes.
Memory monitoring	Memory Usage	memusage	%	Memory usage of Mongos nodes.
Network Monitoring	Private network inbound traffic.	netin	Bytes	Statistics of inbound traffic bytes of Mongos nodes.
	Private network outbound traffic.	netout	Bytes	Statistics of outbound traffic bytes of Mongos nodes.
Latency monitoring	Average latency of all requests	node_avg_all_request_delay	ms	Average latency of all requests received by Mongos nodes.
	Update average latency	node_avg_update_delay	ms	Average latency of update commands on Mongos nodes.
	Insert average latency	replicaset_node	ms	Average latency of insert commands on Mongos nodes.
	Average read latency	node_avg_read_delay	ms	Average latency of read commands on Mongos nodes.
	Average latency of aggregate requests	node_avg_aggregate_delay	ms	Average latency of aggregate commands on Mongos nodes.
	Average latency of Count	node_avg_count_delay	ms	Average latency of counts commands on Mongos nodes.
	Average latency of Getmore	node_avg_getmore_delay	ms	Average latency of Getmore commands on Mongos nodes.
	Average latency of delete	node_avg_delete_delay	ms	Average latency of delete commands on Mongos nodes.
	Average latency of Command	node_avg_command_delay	ms	Average latency of Command on Mongos nodes. Command refers to all commands except insert, update, delete, and query.
	10–50 milliseconds	10ms	Times	The number of requests executed per second with a duration between 10ms and 50ms.
	50–100 milliseconds	50ms	Times	The number of requests per second executed in 50ms to 100ms.
	100ms	100ms	Times	The number of requests per second executed in over 100ms.

©2013–2025 Tencent Cloud. All rights reserved.

Request Monitoring	Total Requests	qps	Times/second	The number of all requests per second on a Mongos node.
	Insert Requests	inserts	Times/second	The number of insert requests per second on a Mongos node. Insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.
	Read Requests	reads	Times/second	The number of read requests per second on a Mongos node.
	Update requests	updates	Times/second	The number of update requests per second on a Mongos node.
	Delete requests	deletes	Times/second	The number of delete requests per second on a Mongos node.
	Count Request	counts	Times/second	The number of count requests received per second on a Mongos node.
	Getmore request	getmores	Times/second	The number of getmore requests received per second on a Mongos node.
	Aggregates requests	aggregates	Times/second	The number of aggregate requests per second on a Mongos node.
	Command requests	commands	Times/second	The number of command requests received per second on a Mongos node. Command refers to all commands except insert, update, delete, and query.
Requests	Total Requests	node_success	Times	Total number of requests received on a Mongos node.
	Number of insert requests	node_inserts	Times	Number of insert requests received on a Mongos node. Insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.
	Number of read requests	node_reads	Times	Number of read requests received on a Mongos node.
	Number of update requests	node_updates	Times	Number of update requests received by a Mongos node.
	Number of delete requests	node_deletes	Times	Number of delete requests received by a Mongos node.
	Number of	node_counts	Tim	Number of count requests received on a Mongos node.

©2013–2025 Tencent Cloud. All rights reserved.

	count requests		es	
	Getmore requests	node_getmores	Times	Number of getmore requests received on a Mongos node.
	Aggregates requests	node_aggregates	Times	Number of aggregate requests received on a Mongos node.
	Command requests	node_commands	Times	Number of command requests received on a Mongos node. Command refers to all commands except insert, update, delete, and query.

Note:
Note: For MongoDB 4.4 (WT.44.13.1) and later versions, insert requests are counted by the actual number of rows inserted. For example, if insertMany() actually inserts 10 records, the monitoring will count 10 insert requests.

Monitoring and Alarms

- For monitoring data and views of each metric, see Viewing Monitoring Data. (Note: Unordered list content)
- For configurable monitoring metrics and how to configure alarms, see Configuring Metric Alarms. (Note: Unordered list content)
- For configurable alarm practices and how to create event rules, see Configuring Event Alarm. (Note: Unordered list content)

©2013–2025 Tencent Cloud. All rights reserved.

Viewing Monitoring Data

Last updated: 2025-02-08 10:33:19

<TencentDB for MongoDB allows you to view the change trend of each monitoring metric, helping you stay up to date with the running status and performance of database resources, make prejudgments, and prevent risks.

Background

- Tencent Cloud Observability Platform (TCOP) is a service that enables real-time monitoring and alarm of cloud service resources. It collects various monitoring metrics data of cloud products and displays them through visual charts, helping you intuitively understand the operation status and performance of cloud products. For more information, see Tencent Cloud Observability Platform.

> **Notes:**
> The original product name of Tencent Cloud Observability Platform was Cloud Monitor. On February 23, 2023, Cloud Monitor was renamed to **"Tencent Cloud Observability Platform"**. For more information, see About the Cloud Monitor Product Name Change.

- In TencentDB for MongoDB, you can use TCOP to create dashboards and various types of charts to compare the metric data of multiple instances. In this way, you can efficiently analyze the changes of monitoring metrics. You can also use TCOP to configure real-time alarms for exceptions during database operations, allowing you to remove risks as soon as they arise.

Version Description

Currently, all versions of MongoDB support instance monitoring.

Billing Instructions

- Basic TCOP features such as alarming and monitoring data collection are free of charge.
- Currently, only **alarm SMS and phone calls** are billed.

Notes

- The monitoring data is retained for 30 days.
- After receiving the alarms reported by Tencent Cloud, you need to troubleshoot problems accordingly.

Prerequisites

- Tencent Cloud Observability Platform service activated.
- You have applied for a TencentDB for MongoDB instance.

Directions

Quickly viewing instance monitoring data

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the **Monitoring/Status** column of the target instance, click to open the instance monitoring panel, where you can quickly view the instance monitoring data.
   - You can view the monitoring data within the corresponding time period by selecting **Real-time**, **Last 24 hours**, **Recent 7 days** or any time period.
   - On the **Requests**, **Connections**, **Capacity and QPS**, or **Latency** tabs, you can view the data of monitoring metrics in different categories.

©2013-2025 Tencent Cloud. All rights reserved.

- In the dropdown list of **Time Granularity**, you can set the collection granularity of monitoring data and obtain fine-grained monitoring data.

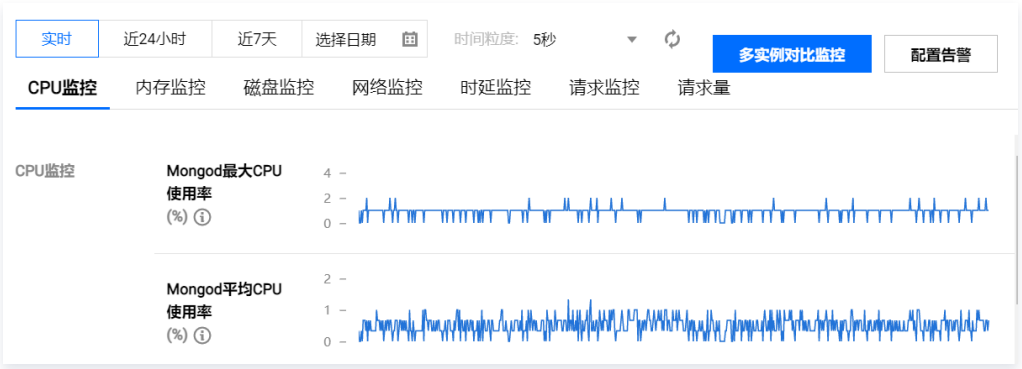

- Click **Multi-Instance Comparison Monitoring** to enter the TCOP **Dashboard List** page. Create a Dashboard, select the relevant instances to be monitored, and set the monitoring chart. You can compare the monitoring data of multiple instances in the same chart as shown below.

- Click **Configure Alarms** to enter the **Create Alarm Policy** page on TCOP. Set **Policy Type** to **TencentDB for MongoDB Instance**, select an **Alarm Object**, and set the **Trigger Condition** for the monitoring metric, configure the alarm notification method. This helps you understand business exceptions promptly, prevent risks, and avoid failures. For detailed directions, see Create Alarm Policy.

Viewing Monitoring Details

1. In the Instance List, find the target instance.
2. Click the Target Instance ID to enter **Instance Details**page.
3. Click the **System Monitoring** tab to view the change trend of each monitoring metric of the entire cluster as shown below (with a replica set instance as an example).

©2013-2025 Tencent Cloud. All rights reserved.

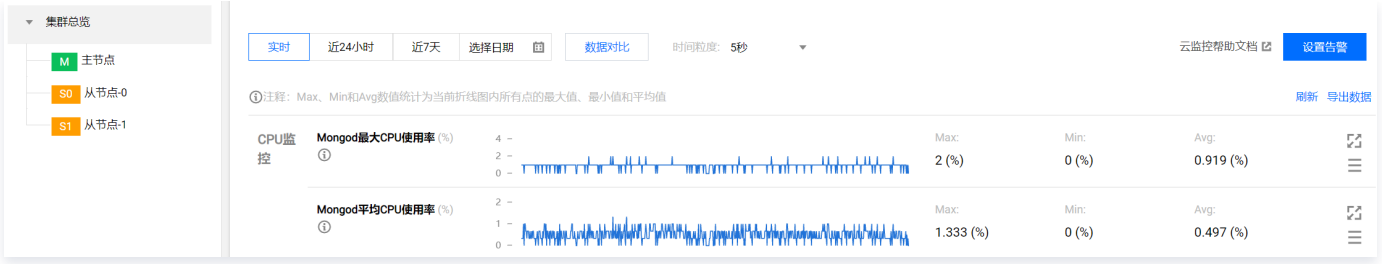

Viewing monitoring data by monitoring object

- Replica set: On the **System Monitoring** page, select the specific instance name, primary node, and secondary node under the **Cluster Overview** to view the monitoring metric data of different monitoring objects.
- Sharded instance: On the **System Monitoring** page, under the **Cluster Overview** cascading navigation node, select the specific shard name, primary node, and secondary node to view the monitoring metrics data of different monitoring objects.

Viewing Monitoring Data by Time Period

On the upper right side of the **System Monitoring** page, you can select **Real-time**, **Last 24 hours**, **Recent 7 days** or any time period to view the monitoring data within the corresponding time period.

Viewing Monitoring Data with Different Time Precision

On the upper right side of the **System Monitoring** page, from the dropdown list after **Time Granularity**, you can select **5 seconds**, **1 minute**, **5 minutes** or **1 day** to view monitoring data with different time precision.

Zooming in on the Change Graph of a Single Metric

In the monitoring metrics list on the right side of the **System Monitoring** page, find the metric you want to view and click . You can zoom in on the metric's change graph, select a time period, set the time granularity, and analyze the metric's changes in more detail.

Export Monitoring Charts

- Export a single monitoring metric chart: In the monitoring metrics list, select the metric to be exported and click . Select **Export Image** to export the metric change graph. Select **Export Data** to view and analyze the monitoring data locally using Excel.
- Export monitoring data in batches: Click **Export Data** above the monitoring metrics list. In the **Export Data** window, select the metrics to be exported and click **Export**. You can then use Excel to view and analyze the monitoring data locally.

Setting Alarms

In the upper right side of the instance monitoring page, click **Set Alarms** to enter the **Create Alarm Policy** page in TCOP. Set the **Policy Type** to **TencentDB for MongoDB Instance**, select an **Alarm Object**, set the **Trigger Condition** of the monitoring metric, and configure the alarm notification method. This helps you stay on top of metric exceptions, prevent risks, and avoid failures promptly. For detailed directions, see Configuring Alarm .

Data Comparison

In the upper right side of the instance monitoring page, you can click **Data Comparison** and set the time range for comparing monitoring data. By default, the data within the past hour is obtained. The curves of the monitoring metric on today and yesterday within the specified time range are displayed in different colors.

©2013-2025 Tencent Cloud. All rights reserved.

Configuring Metric Alarms

Last updated: 2025-02-08 10:33:57

Overview

This is to prevent the system from being affected when some monitoring metrics reach a certain value. You can set alarm rules for these monitoring metrics to enable the alarm system to automatically check the monitoring data and send an alarm notification to the administrator when the monitoring data meets the conditions. This helps you understand business exceptions the first time and resolve them quickly.

Alarm Monitoring Metrics

Alarm Policy Categories

TencentDB for MongoDB provides alarm configurations in three dimensions: instance, replica set, and node. You can set alarm rules for the metrics of each dimension. Among them:

- **Instance**: The instance dimension targets the entire MongoDB cluster, monitoring the number of requests, disk usage, latency, and connections of the entire cluster.
- **Replica Set**: Each replica set in TencentDB for MongoDB follows a master-slave architecture, and each shard in a Sharded Cluster is also a replica set structure, where database documents are stored. This dimension monitors the architecture storing documents, including cache dirty data, cache usage, request hit rate, disk usage, Oplog retention duration, and master-slave delay.
- **Nodes (Mongod, Mongos)**: This dimension targets all nodes in the database cluster, monitoring the usage of Mongod and Mongos nodes, including CPU, memory, disk, inbound and outbound traffic, read and write request counts, queue wait statistics, and connection counts.

Alarm Metrics

Before configuring alarms, please understand the monitoring metrics defined for different policy dimensions and some alarm configuration recommendations for key metrics. For metrics without configuration recommendations, please configure them reasonably based on actual business needs.

Instance Dimensionality

Monitoring Metric Name	Unit	Metric notes	Alarm Configuration Recommendations
Number of write requests	Times	Number of write requests received by the instance.	-
Number of read requests	Times	Number of read requests received by the instance.	-
Number of update requests	Times	Number of update requests received by the instance.	-
Number of delete requests	Times	Number of delete requests received by the instance.	-
Number of count requests	Times	Number of total requests received by the instance.	-
Number of aggregated requests	Times	Number of aggregated requests received by the instance.	-
Number of successful requests	Times	Number of successful requests executed by the instance.	-
Disk Usage Rate	%	Indicates the percentage of the total disk space that is currently used.	The statistical period is 1 minute. If this value is >= 90% and the indicator remains abnormal for 3 consecutive data

©2013-2025 Tencent Cloud. All rights reserved.

			points, an alarm will be triggered once every 30 minutes.
number of delays per unit time (10ms – 50ms)	Times	The number of requests executed between 10 milliseconds and 50 milliseconds.	–
number of delays per unit time (50ms – 100ms)	Times	The number of requests executed between 50ms and 100ms.	–
number of delays per unit time (above 100ms)	Times	The number of requests executed in more than 100ms.	The statistical period is 1 minute. If this value is >= 100 for 3 consecutive data points, an alarm will be triggered once every 30 minutes.
Connection usage rate	%	The percentage of maximum connections currently used in the cluster.	The statistical period is 1 minute. If this value is >= 90% and the indicator remains abnormal for 3 consecutive data points, an alarm will be triggered once every 30 minutes.
Number of requests per second	Times	The number of requests received per second by the instance.	–
Number of command requests	Times	The number of command requests received by the cluster. Command refers to all commands except insert, update, delete, and query.	–
Connection Count	Times	Number of TCP connections for cluster clients.	–

Replica Set Dimension

Chinese Name of Monitoring Indicator	Unit	Metric notes	Alarm Configuration Recommendations
Cache percentage of dirty data	%	The size of dirty data in the cache (bytes) as a percentage of the maximum cache.	The statistical period is 1 minute. If this value is >= 20% for 3 consecutive data points, an alarm will be triggered for the anomaly once every 30 minutes.
Cache usage percentage	%	The ratio between the actual occupied capacity in the cache and the configured maximum cache.	–
Disk Usage Rate	%	Indicates the percentage of the total disk space that is currently used.	The statistical period is 1 minute. If this value is >= 90% and the indicator remains abnormal for 3 consecutive data points, an alarm will be triggered once every 30 minutes.
Cache hit rate	%	Indicates the ratio between the number of requests to retrieve data from the cache and the total number of requests.	–
Oplog retention time	Hour	Oplog is used to record the operation logs of the database, and this metric tracks its retention duration.	–
Average delay between master and	s	In the replica set architecture, a secondary node periodically polls the primary node's	The statistical period is 1 minute. If this value is >= 1800 for 3

©2013–2025 Tencent Cloud. All rights reserved.

slave per unit time		oplogs to replicate the primary node's data. This metric measures the latency of master–slave data synchronization.	consecutive data points, an alarm will be triggered once every 30 minutes.

Mongod node

Chinese Name of Monitoring Indicator	Unit	Metric notes	Alarm Configuration Recommendations
CPU Utilization	%	The percentage of time the CPU is occupied by executing processes to the total CPU time.	The statistical period is 1 minute. If the value is >= 80% for 3 consecutive data points, an anomaly alert will be triggered every 30 minutes.
Memory usage rate	%	Indicates the percentage of space currently used in memory to the total memory capacity.	–
Network Inbound Traffic	MB/s	The amount of inbound traffic per second per node.	–
Network Outbound Traffic	MB/s	The amount of outbound traffic per second per node.	–
Number of waiting read requests in the queue	Count	number of waiting read requests in the queue.	The statistical period is 1 minute. If this value is >= 40 for 3 consecutive data points, an alarm will be triggered for the anomaly once every 30 minutes.
Number of waiting write requests in the queue	Count	number of waiting write requests in the queue.	The statistical period is 1 minute. If this value is >= 40 for 3 consecutive data points, an alarm will be triggered for the anomaly once every 30 minutes.
Connection Count	Count	Number of connected clients.	–
Node Disk Utilization	MB	Utilized capacity of the node disk.	–
ActiveRead of WT engine	Count	Number of read requests for data in memory.	The statistical period is 1 minute. If this value is >= 40 for 3 consecutive data points, an alarm will be triggered for the anomaly once every 30 minutes.
ActiveWrite of WT engine	Count	Number of write requests in memory.	The statistical period is 1 minute. If this value is >= 40 for 3 consecutive data points, an alarm will be triggered for the anomaly once every 30 minutes.
Number of data entries deleted by ttl	Count	Number of data entries automatically deleted by the database after ttl expiration.	–
TTL operation cycles	Times	Number of data checks within the ttl time set in the database.	–

A Mongos Node

Chinese Name of Monitoring Indicator	Unit	Metric notes	Alarm Configuration Recommendations
CPU Utilization	%	The percentage of time the CPU is occupied by executing processes to the total CPU time.	The statistical period is 1 minute. If the value is >= 80% for 3 consecutive data points, an anomaly alert will

©2013–2025 Tencent Cloud. All rights reserved.

			be triggered every 30 minutes.
Memory usage rate	%	The percentage of used space in the memory of the current Mongos node to the total memory capacity.	-
Network Inbound Traffic	MB/s	The amount of inbound traffic per second per node.	-
Network Outbound Traffic	MB/s	The amount of outbound traffic per second per node.	-

Description

- The TCOP configures alarm policies to monitor key metrics of instances, which can be used for free.
- Currently, only **alarm SMS and telephone alarms** are charged. For details, see Billing Overview.

Configuring Metric Alarms

Prerequisites

Activate the TCOP (Tencent Cloud Observability Platform) service.
The database instance is in Running status.
Information about alarm notification objects has been collected, including emails, SMS, and phone calls.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the row of the target instance, enter the Tencent Cloud Observability Platform **Create Alarm Policy** page through any of the following methods.

- Click **Monitoring/Status** column , in the upper right corner of the instance monitoring data panel, click **Configure Alarm**.

- Click the instance ID in blue font to enter the **instance details** page, click the **system monitoring** tab, and click **Setting Alarms**.

6. On the **Create Alarm Policy** page, see the following table to configure the alarm policy. For the basic concept of alarm policy, see Create Alarm Policy.

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Name	Parameter interpretation
Purpose of Policy	Customize the name of an alarm policy for easy identification.
Remarks	Briefly describe the alarm policies for easy identification.
Monitoring Type	Please choose Cloud Product Monitoring.
Policy Type	Set the policy type to **TencentDB / MongoDB / Instance, TencentDB / MongoDB / Mongod Node, TencentDB / MongoDB / Mongos Node** or **TencentDB / MongoDB / Replica Set**.
Project of Policy	Assign a project to the alarm policy. You can quickly filter all alarm policies under the project in the alarm policy list.
Alert Object	• Select **Instance ID**: the alarm policy is bound to the specified database instance. • Select **Instance Grouping**: the alarm policy is bound to the specified database instance group. See Instance Grouping for how to create an instance group. • Select **All Objects**: the alarm policy is bound to all instances for which the current account has permissions. • Select **Tag**: the alarm policy is bound to all instances associated with the current tag key and tag value.
Trigger Conditions	• **Select Template**: Select the template file from the dropdown list and report alarms according to the preset triggering conditions of the template file. For specific configuration, see Configuring Trigger Condition Template. • **Manual Configuration:** Configure the threshold conditions for triggering an alarm for each metric in the **Metric Alarms** area below. Threshold types of Metric Alarms: ○ Select **Static**: Manually set a constant threshold, and send an alarm when the trigger conditions are met.

©2013–2025 Tencent Cloud. All rights reserved.

	○ Select **Dynamic**: Dynamic thresholds are determined based on the threshold boundaries calculated by machine learning algorithms to detect anomalies. For more information, see Create Alarm Policy.
Alarm Notification	You can select a preset or user-customized notification template. Each alarm policy can be bound to at most three notification templates. For details, see Notification Template.

7. Confirm the configuration is correct and click **complete**. For more alarm introduction, see Alarm Introduction.

Relevant API

API Name	API Feature Description
CreateAlarmPolicy	Create TCOP alarm policy

©2013-2025 Tencent Cloud. All rights reserved.

Configuring Event Alarm

Last updated: 2025-02-08 10:34:27

Overview

TencentDB for MongoDB has been integrated into the Tencent Cloud Observability Platform, supporting the reporting of TCOP events. All TCOP events will be automatically delivered to the EventBridge Tencent Cloud service event bus. EventBridge is a secure, stable, and efficient serverless event management platform. An event is a data record of a status change. Event sources need to publish events to EventBridge according to the CloudEvents specification. For more information about the CloudEvents specification, see CloudEvents – Version 1.0.

Event Target

An event rule can have multiple event targets. Before creating an event rule, please plan the event target types. EventBridge currently supports the following **event targets:**

- Message Push (only supports rules in the Tencent Cloud service event bus)
- CLS Log
- SCF
- Ckafka

TencentDB for MongoDB Events

Event Chinese Name	Event Name	Event type	Subordinate Dimension	Recovery Concept Availability	Event Description	Solution and Suggestion
Number of Connections Exceeding Limit	connectionOverlimit	Exception	Instance Dimension	Yes	Instance Connection Count Exceeds Maximum Limit	• Increase the maximum number of connections or restart the instance. For specific operations, refer to Troubleshooting Excessive Connections. • For database performance tuning, refer to Analysis and Solutions for High Connection Utilization.
Data Primary/Secondary Switch	primarywitch	Exception	Instance Dimension	Yes	The primary node of the instance is abnormal, and a switch to the secondary node occurs. This event may be triggered by a physical server failure.	Please confirm whether the instance status is normal.
Disk space is about to run out.	instanceDiskSpaceLow	Exception	Instance Dimension	Yes	Disk space is about to be full, which may cause the instance to become read-only.	Clean up disk space. For specific operations, refer to Solution for High Storage Space Usage.

©2013-2025 Tencent Cloud. All rights reserved.

Node CPU is abnormal.	NodeCPUAbnormal	Exception	Instance Dimension	Yes	An alert is triggered when the CPU utilization of any node in the cluster reaches 80%.	For specific operations, refer to Solution for High CPU Utilization.
SSL Certificate is about to expire	SSLCertAboutToExpire	Exception	Instance Dimension	No	SSL Certificate used by the instance is about to expire	To update the current instance's SSL Certificate, refer to Enabling SSL Authentication.
Node OOM	NodeOom	Exception	Instance Dimension	No	Mongod node memory usage overload	Evaluate whether the current database memory specifications meet business needs. If additional memory is required, consider upgrading the configuration. For details, see Adjust the mongod node specification.
Node Restart	NodeReboot	Exception	Instance Dimension	No	Mongod node restart	If it is an unplanned reboot (non-manual operation reboot, parameter change, instance upgrade, node migration, etc.), check whether the business load is too large, and consider upgrading the configuration if necessary.
Node status is abnormal	NodeAbnormal	Exception	Instance Dimension	Yes	Mongod or Mongos Node status exception, may affect service	Manual inspection of node status. If the issue cannot be resolved, please submit a ticket.

Billing Instructions

Tencent Cloud offers EventBridge **pay-as-you-go** purchase method. For more information, see EventBridge > Product Pricing.

Type	Pay-As-You-Go
Payment Mode	Hourly settlement based on the number of events actually delivered to the event set
Billing Unit	CNY/million events
<0>Application cases<0>	In application scenarios with low message volume or high message volume fluctuation, resource waste can be effectively avoided

Directions

1. Log in to the Tencent Cloud EventBridge Console and select Event Rule in the left sidebar.
2. At the top of the right page, select **Region** as **Guangzhou**, and choose **default** from the **Event Set** dropdown list.

> **Notes:**
> - Tencent Cloud service event bus is used to collect monitoring events and auditing events generated by Tencent Cloud services in all regions. It is created in Guangzhou by default and cannot be deleted.
> - In the left sidebar, select **Event Bus**. In the event set list, click **default** to see that the default **default** event set already includes TencentDB for MongoDB. For specific operations, see Tencent Cloud Service Event Source.

3. On the **Event Rule** page, click **Creation**. In the **Event Mode** navigation page, configure the page parameters according to the parameter explanation table below.

©2013-2025 Tencent Cloud. All rights reserved.

Interface Area	Interface Parameter	Parameter interpretation
Basic Information	Region	Create the event rule region.
	Event Bus	Event set information to which the event rule belongs.
	Rule Name	Set the event rule name. It can contain 2-60 letters, digits, underscores, and hyphens and must start with a letter and end with a digit or letter.
	Rule Description	Provide a brief description of the event rule.
	Tag	Set the tag key and value for the event.
	Data conversion	Check if data conversion is needed.
Event Example	Event Example Selection	In the dropdown list, you can search for MongoDB to view examples of MongoDB events.
Event Match	Write Mode	• **Form Mode**: In this mode, you can select **Cloud Service Type** and **Event Type**, and provide event matching rules. • **Custom Event**: In this mode, please enter custom event matching rules in the input box below. For how to write rules, please click **Rule Writing Guide**.
	Cloud Service Type	When **Write Mode** is set to **Form Mode**, this parameter is displayed. In the dropdown list, select **TencentDB for MongoDB**.
	Event Type	When **Write Mode** is set to **Form Mode**, this parameter is displayed. In the dropdown list, select the supported event type.
	Event Matching Rule Preview	Preview the generated event matching rules.

4. Click **test matching rules** to test the defined event matching rules. After passing the test, click **Next**. If the test fails, correct it according to the prompt information.
5. (Optional) If you need to convert the data format, the **Event Transformation** page will be displayed as shown below. Configure the data transformation format and fields according to the parameter explanation in the table below.

> **Notes:**
> Data transformation provides a simple data processing feature. By passing in data and configuration items, it can format the data, return the processed structured data, and distribute it to downstream targets, bridging data sources and data processing systems.

©2013–2025 Tencent Cloud. All rights reserved.

Interface Area	Interface Parameter	Parameter interpretation
Create New Data Transformation	Event Pattern Preview	• Select **sample event** to use the event template. • Select **manual input** to customize event fields in the input box below.
	Event Template	**Event Pattern Preview** Select **Sample Event** to display this parameter. In the dropdown list, you can search for MongoDB and select the MongoDB event template. The specific field information of the event template will be displayed in the input box below.
	Transformation Target	• **Complete Event**: Routes the complete structure of event fields to the event target. • **Partial Event**: EventBridge extracts event fields configured by **JSONPath** and routes the specified event fields to the event target.
	JSONPath	When selecting partial event for transformation target, this parameter is displayed. Please enter the event fields to be transformed in the input box.
	Parsing Mode	Select the parsing mode, supporting JSON, Separator, and Regex.
	Resolved Result	Click **OK** behind the parsing mode to start parsing data and convert event rules into Key-Value format.
	Filter	Configure the filter to output only data that meets filter rules.
	Data Processing	For the currently parsed data, please select the data type in the **TYPE** column.
	Test Results	Click **Testing** to perform a legality check and output the final converted results.
Failure Information Handling	Dead Letter Queue	Configure whether to deliver messages that cannot be processed normally to the dead-letter queue of Ckafka.
	Delivery Type	The delivery type of fixed failure messages is Ckafka.
	CKafka	Select the instance ID of Ckafka for failed message delivery.

©2013-2025 Tencent Cloud. All rights reserved.

	instance	
	CKafka Topic	Select the Topic of the Ckafka instance for failed message delivery. CKafka uses the concept of Topic externally, where producers write messages to the Topic and consumers read messages from the Topic.

6. Click **OK** next to the parsing mode to start parsing data. Wait for data parsing to complete, then set filter rules and data processing method. For detailed operations, see Configuring Data Conversion.
7. Click **Next**, select the event target bound to this rule, and you can deliver the collected events to the specified delivery target for complete processing and consumption. The following figure takes **Trigger method** as **Notification message** as an example. For configuring event alert push, refer to Configure delivery targets.

8. For the event rule to take effect immediately, check **Enable Event Rule Immediately** and click **Done**.

Event Rule related interfaces

API Name	API Feature
CheckRule	Validate rules
CreateRule	Creating Event Rule
DeleteRule	Delete Event Rule
GetRule	Get Event Rule details
ListRules	Get Event Rule list
UpdateRule	Update Event Rule

More Operations

To view, edit, or delete an event rule, see Managing Event Rule.

FAQs

For event rule related concepts and billing related frequently asked questions, see Event Bus > FAQs.

©2013–2025 Tencent Cloud. All rights reserved.

Backup and Restore
Data Backup

Last updated: 2025-02-08 10:35:10

To prevent data loss due to system failures and other factors, TencentDB for MongoDB supports data backup and data rollback after system recovery to ensure data integrity.

Backup Type

- Automatic backup: refers to the scheduled automatic backup of data described in the system's default backup policy (such as the default backup time interval and backup method).
- Manual backup: refers to immediate execution of backup tasks based on business ops troubleshooting needs.

Backup method

In MongoDB, backup strategies are usually divided into logical backup, physical backup, and snapshot backup.

> **Notes:**
> - General version instances support logical backup and physical backup.
> - Cloud disk version instances support physical backup and snapshot backup, but do not support logical backup currently.

- **Physical backup**: It refers to directly copying the related physical files of the database to another location. This method backs up the underlying storage files of the database, including data files, index files, and log files. The backup speed is fast, as it directly copies files without parsing the database content, resulting in high backup and rollback efficiency.
- **Logical backup**: It uses the mongodump tool to store the operation logs of the database into logical backup files for data backup. During recovery, the data is restored by replaying the operation logs. This method backs up the logical structure and data of the database, including documents, indexes, collections, and database-level configurations. The backup speed is slow, but it has strong portability, allowing the logical backup of the database to be restored to different versions of the database. However, it may require more storage space and could impact database performance during the backup process.
- **Snapshot backup**: It is a data backup method that creates a snapshot of the cloud disk at the storage layer, capturing the disk data state at a specific point in time. The backup speed is fast, with minimal impact on database performance, and no need for downtime. Data can be easily restored from the snapshot.

Use Limits

The backup can continuously cover data for 7 days, allowing rollback to any time point within the last 7 days.

Must-Knows

- Business usage is not affected during instance backup.
- Backup files are stored in Tencent Cloud Object Storage (COS), which does not occupy the storage space of TencentDB for MongoDB instances. For more information about the object storage service, see COS.

> **Notes:**
> - During the instance creation process, an automatic backup will be triggered based on the instance's default backup policy.
> - After scaling specifications or cross-AZ migration, a new automatic backup will be triggered again.

Version Description

General Edition

Automatic and manual backups for General Edition instances in versions 3.2 and 3.6 only support logical backup, while versions 4.0 and later support physical backup.

Version	Automatic Backup	Manual Backup

©2013-2025 Tencent Cloud. All rights reserved.

Versions 3.2 and 3.6	Supported backup method: logical backup	Supported backup method: logical backup
Version 4.0 and above (4.0, 4.2, 4.4, 5.0, 6.0, 7.0)	Default backup method: logical backup Supported backup methods: logical backup and physical backup	Default backup method: logical backup Supported backup methods: logical backup and physical backup

Cloud Disk Edition

Cloud disk version instances support versions 4.0, 4.2, 4.4, and 5.0, and support snapshot backup and physical backup.

Version	Automatic Backup	Manual Backup
4.0, 4.2, 4.4, 5.0	Default backup method: snapshot backup Supported backup methods: snapshot backup and physical backup	Default backup method: physical backup Supported backup methods: snapshot backup and physical backup

Description

Currently free of charge, future backup storage fees will be notified separately.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Adjusting the automatic backup policy

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. Select the **Backup and Restore** tab to enter the **Backup Task List** page.
7. Select the **Automatic Backup Settings** tab and click **Edit**.

> **Notes:**
> The following figure shows a cloud disk version instance. The backup method does not support logical backup currently.

8. Re-edit the **backup method** and **backup start time** according to the parameter description in the table below.

Parameters	Description
Data Backup Retention	By default, backup data is retained for **7** days.
Backup method	(Optional) Select a backup method, including: logical backup, physical backup, snapshot backup. • How to choose a backup strategy: Please refer to the description of Backup Mode to select a strategy that suits your needs. • The supported backup methods vary by version. For specific information, see Version Notes.

©2013-2025 Tencent Cloud. All rights reserved.

	Notes: MongoDB Community Edition 3.6 replica set instance does not support setting this parameter.
Backup Start Time	The default start time is 22:00–02:00, meaning the system will start the backup task between 22:00–02:00 every day. • Supports selecting different time periods to start backing up data, which can be set according to actual business situations. • The specific start time will vary depending on the specific scheduling of the backup task.

9. Click **Save** to save the settings. Subsequent backups will be performed automatically according to the backup rules.

Manual Backups

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. In the top right corner of the **Instance Details** page, click **Manual Backup**.
7. (Optional) In the pop-up dialog box, select **Backup Method**. This parameter is not supported for MongoDB 3.6 replica set instances.
8. Add remark information, click **Yes**, and the data is saved for 7 days by default.

Downloading Backup Files

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. Select the **Backup and Restore** tab to enter the **Backup Task List** page.
7. In the **Backup Task List**, find the file to be backed up, and click **Download** in the **Operation** column.
8. In the **Generate Backup File** dialog box, read the backup prompt information, and click **Yes**.
9. Click the **Download File List** tab to view the backup task progress.
10. After the task is completed, back up the data to the local machine as follows.
    - External network method: In the **Operation** column, click **Public Network Download** to directly download the backup to the local machine via the browser.
    - Internal network method: Copy the intranet address, and on the CVM server, use the wget command format: `wget -c 'intranet address' -O backup.tar` for high-speed internal network download. For how to log in to CVM, see Log in to CVM.

Relevant API

API	Description
DescribeDBBackups	Query instance backup list
CreateBackupDBInstance	Back up an instance
DescribeBackupDownloadTask	Query backup download task information
CreateBackupDownloadTask	Create backup download task

©2013–2025 Tencent Cloud. All rights reserved.

Roll back data
Clone data

Last updated: 2025-02-08 10:35:47

Overview

When the data of the current instance encounters severe issues and needs to be rolled back to a previous backup state, you can quickly restore data by directly cloning a new instance from the backup file of the current instance. The data of the cloned instance is consistent with the backup file. You can use the cloned instance to analyze historical data or achieve the purpose of rollback by swapping the IPs of the cloned new instance and the original instance. This method can avoid the tedious process of manually restoring data one by one and improve the efficiency and accuracy of data recovery.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharded instance is in Running status.
- Data backup completed.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select Replica Set Instance or Sharded Instance. The directions for the two types of instances are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click **Target Instance ID** to enter the instance details page.
6. Select the **Backup and Restore** tab to enter the Backup Task List page.
7. In the Backup Task List, find the backup file to be restored.
8. In its **Operation** column, click **Clone**.

9. On the **Clone TencentDB for MongoDB Instance** page, confirm the master instance information, select the restore time point in the time frame of **Select Restore Time Point**, choose the billing mode and configuration specification of the new instance, and purchase the new instance. For more information on how to configure parameters, see Creating TencentDB for MongoDB Instance.

> **Note:**
> The rollback time only supports selecting data from any point in time within the last 7 days before the current time. (Note)

©2013-2025 Tencent Cloud. All rights reserved.

10. Confirm the cost, then click **Buy Now**.
11. Return to the instance list page. Once the instance creation is completed, the data of the source instance is synchronized to the newly cloned instance, and you can use the new instance. You can achieve the purpose of data rollback by changing the IP to swap the IP of the cloned new instance with the IP of the original instance.

> Description
> After the instance is cloned, the source instance can be retained or terminated based on your needs. (Note)

More Entries

1. In the Backup Task List on the **Backup and Restore** tab, find the backup file to be restored.
2. In the **Operation** column, click **Database table rollback**.
3. In the **Batch rollback database table data** configuration wizard, on the **Select rollback instance** tab, select **rollback target instance** as **rollback to new instance**, and select **select rollback type** as **full instance rollback**.
4. In the instance list below, check an instance to be restored (you can search by instance id, instance name, or ip addresses in the search box).

©2013–2025 Tencent Cloud. All rights reserved.

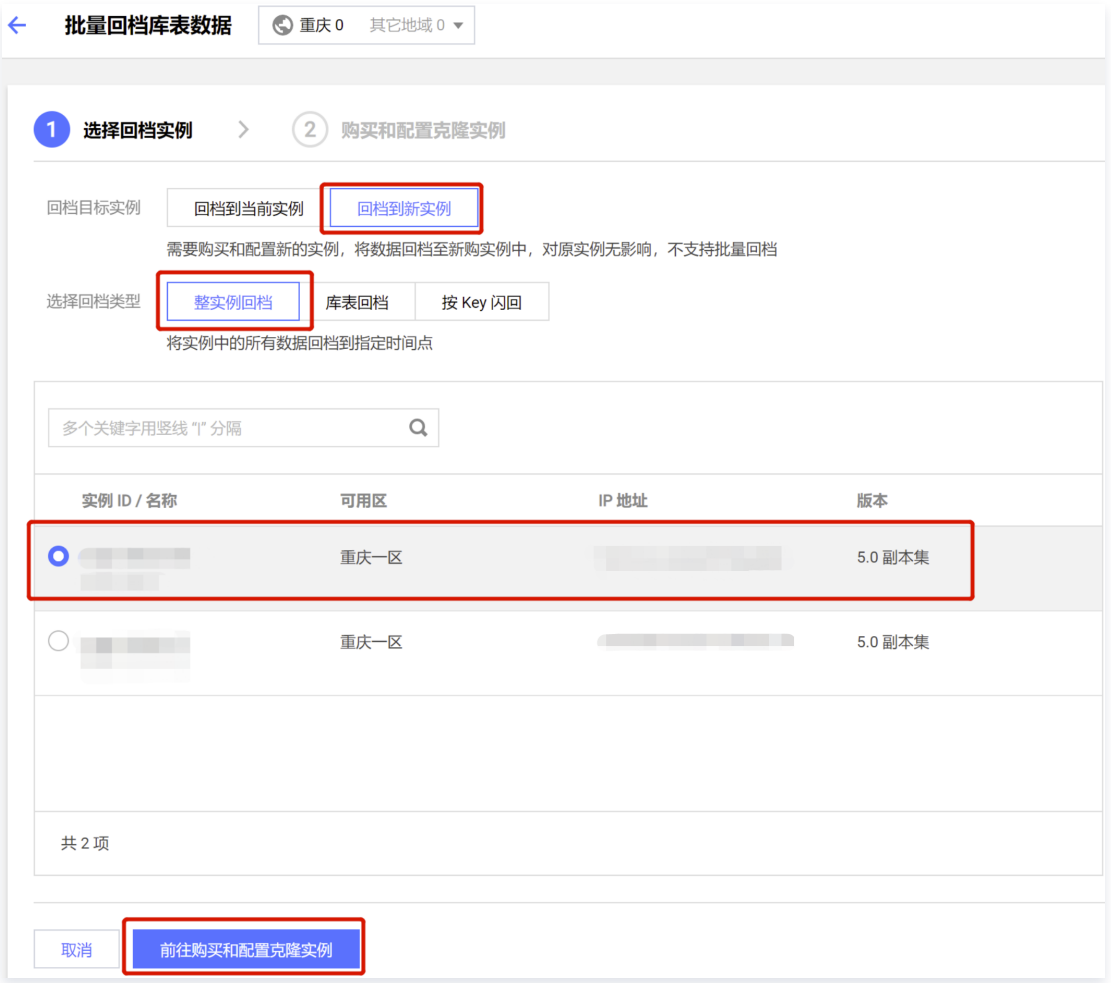

5. Click **Go to purchase and configure clone instance** to enter the **Clone TencentDB for MongoDB Instance** page, confirm the master instance information, select the billing mode and configuration specification of the new instance, and purchase new instance. For more information on how to configure parameters, see Creating TencentDB for MongoDB Instance.
6. Confirm the cost, then click **Buy Now**.

©2013–2025 Tencent Cloud. All rights reserved.

Rolling Back Tables

Last updated: 2025-02-08 10:36:49

Overview

When the business only needs to perform data recovery on multiple tables in the database, you can perform **table rollback** on the console, restoring data in the current instance or a new instance. Compared to full instance rollback, less data is rolled back, making table rollback faster than full instance rollback.

Version Description

Version Information	Rollback methods
3.2, 3.6	• Full Instance Rollback (Logical Backup) • Table Rollback (Logical Backup)
4.0, 4.2, 4.4	• Full Instance Rollback (Logical Backup, Physical Backup) • Table Rollback (Logical Backup, Physical Backup)
5.0	• Full Instance Rollback (Logical Backup, Physical Backup) • Table Rollback (Logical Backup, Physical Backup) • flashback by key
6.0	• Full Instance Rollback (Logical Backup, Physical Backup) • Table Rollback (Logical Backup, Physical Backup)

Use Limits

- You can select up to 2,000 tables per instance to roll back.
- Only supports rolling back data to any time point in the last 7 days.

> ⚠ **Note:**
> - Pay close attention to the **System Monitoring** metric **Oplog Time Difference** on the instance management page. The smaller the metric value, the greater the risk of oplog being overwritten when your business has frequent write, update, and delete operations.
> - Note: If there are transaction operations on the client side during rollback, please actively commit the transactions or set a timeout to avoid long-term lock resource occupation causing rollback task exceptions.
> - Note: After the number of shards in a sharded cluster changes, previous backups cannot be rolled back. Please select a backup made after the shard change for rollback.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB instance is in **Running** status.
- Data backup completed.

Directions

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. On the **Instance Details** page, select the **Backup and Rollback** tab.
6. On the **Backup and Rollback** tab, go to the **Backup Task List** page.

©2013-2025 Tencent Cloud. All rights reserved.

7. In the **Backup Task List**, find the backup file to be restored.
8. In the **Operation** column, click **Database table rollback**.

9. In the **batch rollback database table data** configuration wizard, on the **Select Archived Instance** tab, after **rollback target instance**, choose **Rollback to Current Instance** or **Rollback to New Instance, Select Rollback Type** and select **Table Rollback**.

> **Note:**
> - **Restore to Current Instance**, no need to purchase a new instance, restore table to current instance, support selecting multiple instances for batch restore, choose based on actual scenario for table restore and key-based rollback. In the instance list below, check one or more instances to be restored (you can search by instance ID, instance name, or IP addresses in the search box). For specific operations, see Table Restore to Current Instance.
> - **Rollback to new instance** requires purchasing a new instance, has no impact on the source instance, and does not support batch rollback for multiple instances. You can choose to restore table, flashback by key, or clone instance based on the actual scenario. In the instance list below, you can only select one instance to be restored (you can search by instance ID, instance name, or IP address in the search box). For detailed operations, see Restore Table to New Instance.

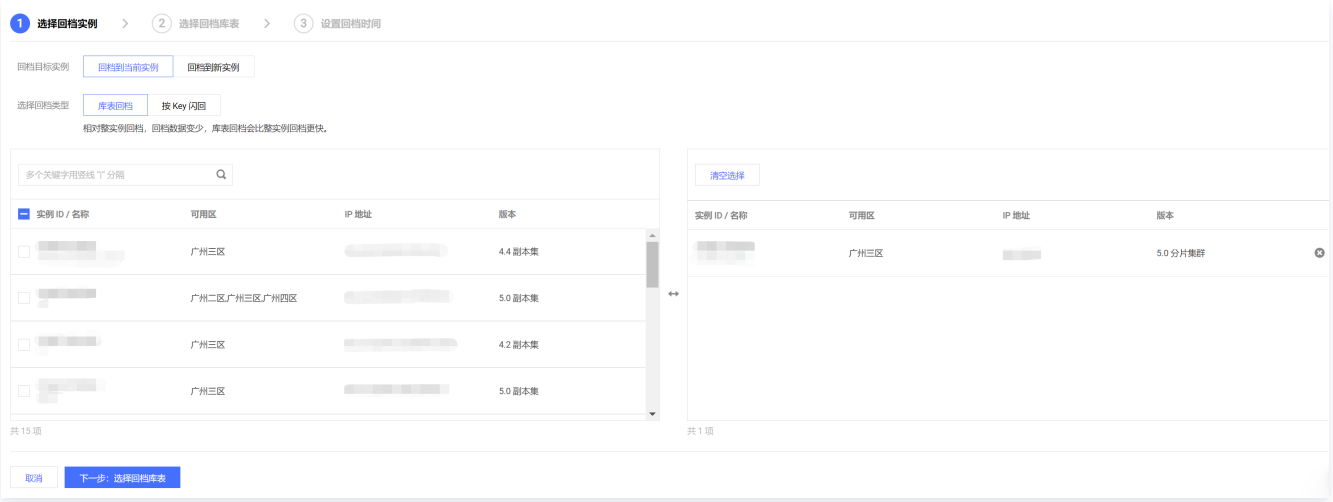

Restoring Tables to Current Instance

1. Click **Next: Choose the Table to Rollback**, on the **Choose the Table to Rollback** tab, select the tables to be rolled back, and confirm the database table information in the right box area. As shown below.

> **Note:**
> In the right area, you can confirm and modify the selected database tables.
> - Click **Clear Selection** to clear the selected tables in case of a selection error.
> - Click ⊗ to delete the selected tables one by one.

©2013-2025 Tencent Cloud. All rights reserved.

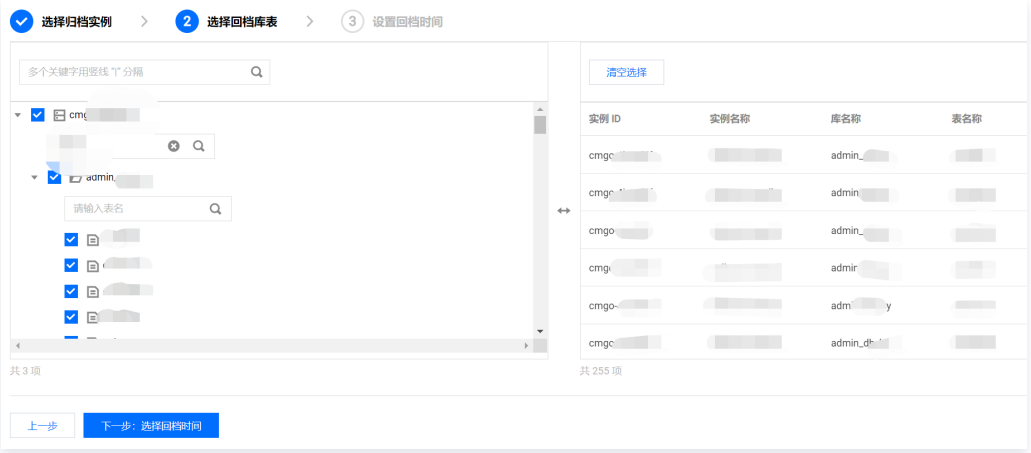

2. Click **Next: Select Rollback Time**, in the **Set Rollback Time** tab, select the specific time point to be rolled back in the time box behind **Set Rollback Time**, and confirm the pre-rollback instance information and database table information.

> **Note:**
> - The rollback time only supports selecting data from any point in time within the last 7 days before the current time. (Note)
> - Restoring to the current instance does not directly restore the original table but creates a new backup file. For example, if the source database table is test, a new test_bak database table will be created. As shown in the figure below, the **rollback table name** is the name of the new database table.
> - After the rollback task is completed, you can batch modify the database table names as needed.
> - For sharded clusters of TencentDB for MongoDB below version 5.0, when rolling back databases and tables to the current instance, the collection name cannot be modified and data can only be manually replaced back to the original collection.
> - To prevent the production database table from being accidentally deleted, if the restore oplog contains drop database or drop collection operations, the rollback task will be terminated.

3. Click **Initiate Rollback** to return to the **Batch Rollback Database Table Data** tab of the **Rollback Task** page, where you can see the ongoing rollback tasks. Click **Task Details** in the **Operation** column to view detailed information about the task. Wait for the task to complete, then connect to the instance to confirm rollback data accuracy.

4. (Optional) Select the **Batch Rollback Database Table Data** tab of **Batch Rename Tables**, find the completed rollback task, and click **Bulk Modify Table Names** in the **Operation** column. You will see the information of tables to be modified in the right area, including the original table name, original table new name, rollback table name, and rollback table new name.

©2013-2025 Tencent Cloud. All rights reserved.

- Click to download the information of tables to be modified for local viewing.
- Confirm the modification, and click **Batch Rename Tables** at the bottom left to complete the modification.

> **Note:**
> - Batch renaming tables can only modify all tables under a single instance in a single rollback task. If the user initiates a batch rollback task that rolls back multiple instances' tables, the table names need to be modified one by one. For specific operations, please refer to Batch Rollback.
> - Batch renaming tables includes modifying the original table name and the rollback table name.
>   - Original table: add the _ori identifier to its original table name.
>   - Rollback table: change the rollback table name to the original table name.

Rolling Back Database Tables to a New Instance

1. Click **Next: Select Rollback Table**, and on the **Select Rollback Table** tab, choose the source instance's tables to be rolled back. In the search box, you can search for the tables to be rolled back by library name and table name. View the selected table information in the right box area as shown below. In the right box area, you can manage the selected tables.
   - Click **Clear Selection** to clear the selected tables if there is a selection error.
   - Click to delete the selected tables one by one.

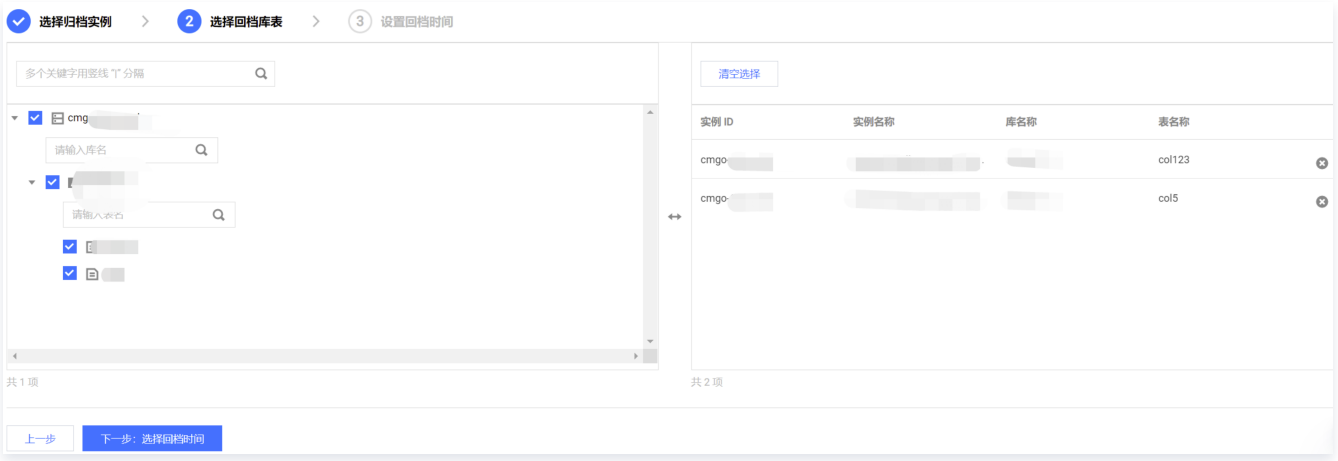

2. Click **Next: Select Rollback Time**, and on the **Set Rollback Time** tab, select the rollback time in the **Set Rollback Time** time box, and confirm the pre-rollback instance information and table information.
3. Click **Go to Purchase and Configure Clone Instance** to enter the **Clone TencentDB for MongoDB Instance** page, and select the billing mode and configuration specification of the new instance. For more information, see Creating TencentDB for MongoDB Instance.
4. Confirm the cost, click **Buy Now**.
5. Return to the instance list page. Once the instance is created, the source instance's tables will be synchronized to the newly purchased clone instance. You can connect to the new instance to confirm the accuracy of rollback data.

©2013–2025 Tencent Cloud. All rights reserved.

flashback by key
Initiate Flashback by Key Task

Last updated: 2025-02-08 10:37:41

Overview

TencentDB for MongoDB supports the feature of flashback by key, enabling real-time backup of specified collections. This facilitates rapid rollback of minor data disorder or accidental loss caused by system vulnerability or failure, based on the flashback key of the data (default is id), ensuring rapid business recovery.

> **Notes:**
> Flashback by Key officially started public beta on September 11, 2023. Please submit a ticket to apply. The Tencent Cloud Database team will invite you for product experience and testing. No additional costs during public beta.

Constraints and limitations

- Currently, only version 5.0 supports the flashback by key feature.
- Supports enabling the flashback by key feature for multiple collections within an instance, with a maximum of 100 collections per instance.

Operation Steps

Step 1: Logging into the Backup and Rollback Tab

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. On the right side, select a **Region** at the top of the Instance List page.
4. In the Instance List, find the target instance.
5. Click the **Target Instance ID** to enter the **Instance Details** page.
6. Select the **Backup and Rollback** tab to enter the default **Backup Task List** page.

Step 2: Enabling the Flashback by Key Feature

1. Select the **key-based rollback** tab and click .
2. In the **key-based rollback settings** configuration wizard, on the **select database tables** tab, select one or more specific databases and tables as shown below, and click **Next**.

> **Notes:**
> In the right area, you can confirm and modify the selected database tables.
> - Click **Clear Selection** to clear the selected tables in case of a selection error.
> - Click to delete the selected tables one by one.

©2013-2025 Tencent Cloud. All rights reserved.

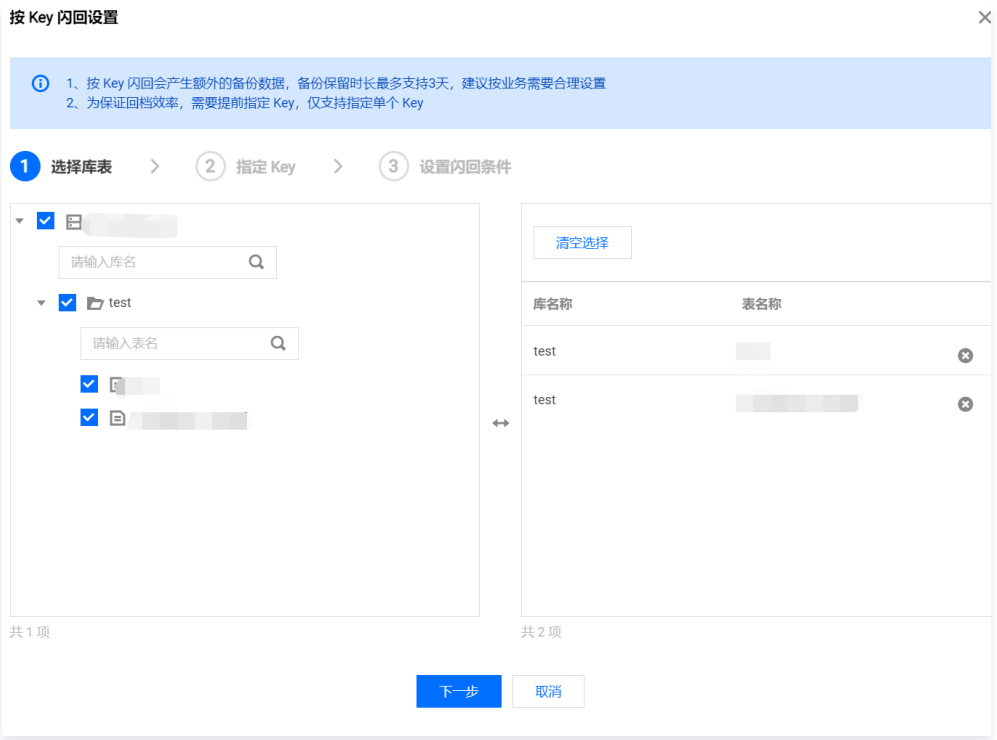

3. In the **key-based rollback settings** configuration wizard, on the **specify key** tab, specify an **additional flashback key** for the selected tables, and click **Next**.

> **Notes:**
> - By default, key-based rollback storage is backed up based on _id and timestamp.
> - An additional flashback key is a specified key used to index the flashback storage. During data flashback, this specified key can be used to quickly retrieve the collection to be restored. If no additional flashback key is specified, the flashback will be based on the default "_id" key.

4. In the **key-based rollback settings** configuration wizard, on the **set rollback policy** tab, set the **backup retention period**.

> **Notes:**
> - The backup retention period refers to the duration for which the backup files generated by the key-based flashback feature can be retained. In TencentDB for MongoDB, key-based flashback currently supports three retention periods: 12 hours, 1 day, and 3 days.
> - When you enable the key-based flashback feature, the system will generate a backup file, which can be retained for the specified period. For example, if you choose a retention period of 1 day, the system will start timing after

©2013-2025 Tencent Cloud. All rights reserved.

generating the backup file and automatically delete the backup file after 1 day.

5. Click **Yes,** return to the key-based flashback tab, and you will see the following interface. The key-based flashback settings are now enabled. Click **Modify** to change the previous configuration.

Notes:

- For the backup retention period, click **Modify** to readjust the retention period.
- Click the **add collection** behind the supported database and table for flashback by key to add the collection to be rolled back.
- In the selected database and table for flashback by key, click **delete** in the operation column to delete the selected collection and re-add the collection.
- For the selected collection, modifying the specified key is not supported.

Step 3: Initiating a Flashback Task

1. On the **Backup and Restore** page, switch to the **Backup Task List** tab.
2. In the **Backup Task List**, find the backup file to be restored.
3. In the **Operation** column, click **Database table rollback**.

4. In the **batch rollback database table data** configuration wizard, on the **select archive instance** tab, after **rollback target instance**, select **restore to current instance** or **rollback to new instance, select rollback type** and choose **key-based rollback**.

Notes:

- **Restore to current instance**, no need to purchase a new instance, restore table to current instance, support selecting multiple instances for batch rollback, and choose table restore and key-based rollback based on actual scenarios.
- **Rollback to new instance**, need to specify another instance. You need to prepare a new instance in advance before the operation, no impact on the source instance, and does not support selecting multiple instances for batch rollback. Choose table restore, key-based rollback, or clone instance based on actual scenarios.

Key-based rollback to the current instance

©2013-2025 Tencent Cloud. All rights reserved.

1. Click **Next: Select Rollback Table**, and on the **Select Rollback Table** tab, choose the source instance's tables to be rolled back. In the search box, you can search for the tables to be rolled back by library name and table name. View the selected table information in the right box area as shown below. In the right box area, you can manage the selected tables.
   - Click **Clear Selection** to clear the selected tables in case of a selection error.
   - Click ⊗ to delete the selected tables one by one.

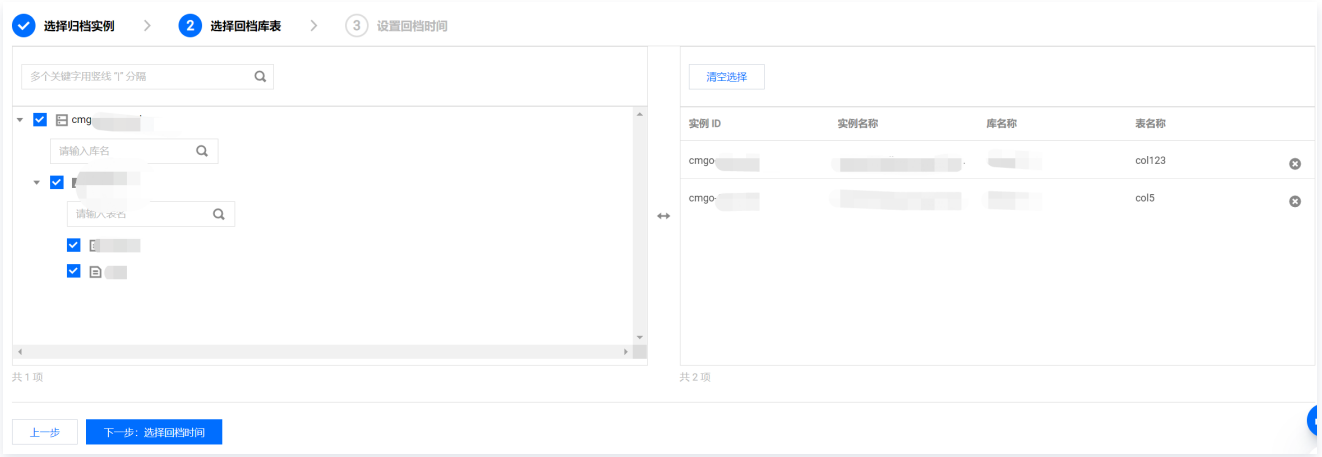

2. Click **Next: Select Rollback Time**, and on the **Set Rollback Time** tab, select the rollback time in the **Set Rollback Time** time box, and confirm the pre-rollback instance information and table information.
   - In the **specify key** dropdown list, select the key to filter documents during rollback. If no additional rollback key is specified when enabling rollback, the default `_id` is used to filter tables.
   - In the **rollback list**, click **input condition value** to specify the key value corresponding to the Key of the flashback document, or click **please upload a file** to upload the key value in a CSV file.

> **Notes:**
> The format and limitations of the input condition value are as follows. Click **Example** to refer to the sample input condition value.
> - The input content only needs to be entered in text form to represent the value of the flashback Key. The system will automatically concatenate it into the `{"key":"value"}` format.
> - Data types need to distinguish between numeric data type and string type. String type is enclosed in double quotes. For example: 20, "hello", "10". If special characters are included, they need to be escaped.
> - Input condition value should be one record per line, with a maximum of 100 lines. If more than 100 lines, click **please upload a file** to upload a CSV file for entry.
>   - The maximum file size for CSV format is 20M.
>   - The rollback list of a CSV file for a collection can have up to 50,000 rows. Exceeding this limit will result in an error.
>   - If there are illegal rows, they will be displayed after the CSV file is uploaded. You need to confirm the illegal information and only rollback the legal records.

3. Click **initiate rollback** and confirm the instance information in the instance rollback pop-up window.

©2013-2025 Tencent Cloud. All rights reserved.

4. Click **Yes** to enter the **batch rollback database table data** task page and wait for the task to complete.

Flashing back by key to a specified instance

1. Click **Next: Select Rollback Table**, and on the **Select Rollback Table** tab, choose the tables of the source instance to rollback. In the search box, you can search for the tables to rollback by database name and table name. View the selected table information in the right box area, and manage the selected tables.
   - Click **Clear Selection** to clear the selected tables in case of a selection error.
   - Click to delete the selected tables one by one.
2. Click **Next: Select Rollback Time**, and on the **Set Rollback Time** tab, select the rollback time in the **Set Rollback Time** time box, and confirm the pre-rollback instance information and table information.
   - In the **specify key** dropdown list, choose the key for filtering documents during the rollback. If no additional rollback key is specified when initiating the rollback, the default `_id` will be used.
   - In the **rollback list**, click **input condition value** to specify the key value corresponding to the Key of the flashback document, or click **please upload a file** to upload the key value in a CSV file.

     Notes:
     The format and limitations of the input condition value are as follows. Click **Example** to refer to the sample input condition value.
     - The input content only needs to be entered in text form to represent the value of the flashback Key. The system will automatically concatenate it into the `{"key":"value"}` format.
     - Data types need to distinguish between numeric data type and string type. String type is enclosed in double quotes. For example: 20, "hello", "10". If special characters are included, they need to be escaped.
     - Input condition value should be one record per line, with a maximum of 100 lines. If more than 100 lines, click **please upload a file** to upload a CSV file for entry.
       - The maximum file size for CSV format is 20M.
       - The rollback list of a CSV file for a collection can have up to 50,000 rows. Exceeding this limit will result in an error.
       - ``
         ```
         Note: If there are illegal rows, illegal records will be displayed after the CSV file is uploaded.
         You need to confirm the illegal information and only rollback legal records.
         ```
         ``
   - In the **rollback to specified instance** dropdown list, specify the target instance to rollback.

     Notes:

©2013-2025 Tencent Cloud. All rights reserved.

```
Note: To avoid potential issues caused by version differences with the original cluster, it is recommended that the target instance for rollback be of the same version as the original collection.
```

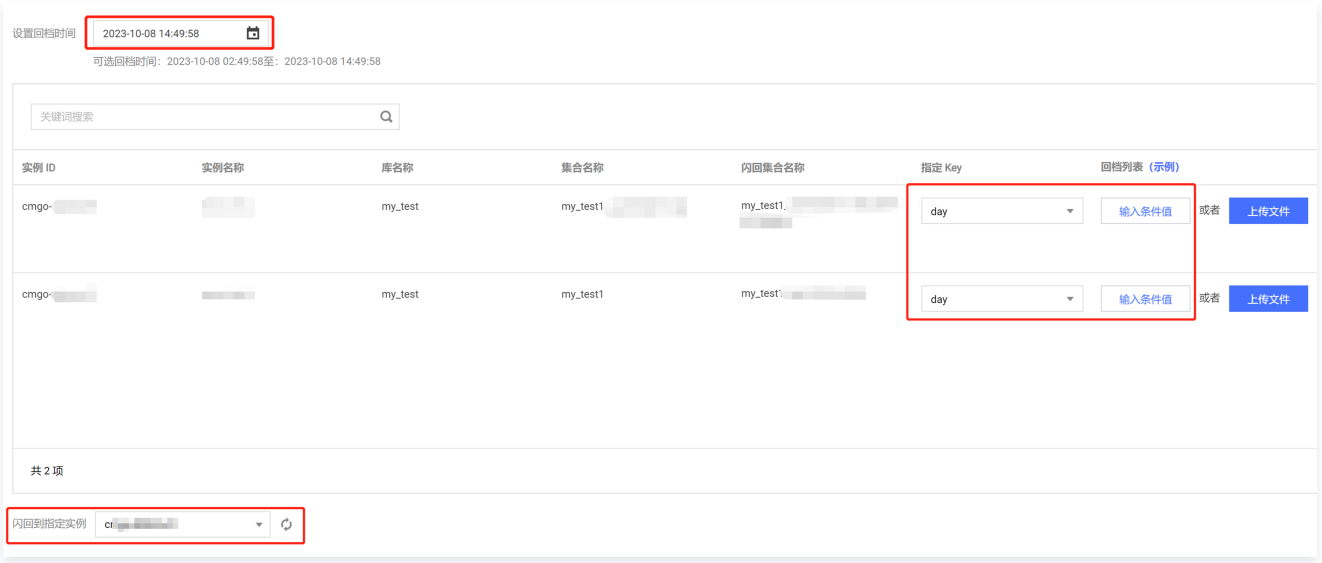

3. Click **initiate rollback** and confirm the instance information in the instance rollback pop-up window.

4. Click **Yes** to enter the batch rollback task page and wait for the task to complete.

Step 4: Batch Updating Source Collection Data

1. On the **batch rollback database table data** task page, wait for the task to complete.

2. In the **Operation** column of the rollback task, click **batch update data**. In the **batch update data guide** window, follow the steps to batch update the original collection data. For specific update examples, refer to the batch update data example.

> **Note:**
> Restoring to the current instance does not directly restore the original table but creates a new backup file. For example, if the source database table is test, a new test_bak database table will be created. As shown in the figure below, the

©2013–2025 Tencent Cloud. All rights reserved.

rollback table name is the name of the new database table.

©2013–2025 Tencent Cloud. All rights reserved.

Batch Update Data Example

Last updated: 2025-02-08 10:38:07

This article explains the data update operation process based on the timeline shown below and the key-based rollback.

Data Preparations

1. If the original collection name is `test`, at time point A, there are three pieces of data in the table, recording three people's scores (score), all of which are 90 points.

```
[direct: mongos] tset> db.test.find()
[
  { _id: ObjectId("652d2a6505726f70625ce5cf"), name: '张三', score: 90 },
  { _id: ObjectId("652d2a6b05726f70625ce5d0"), name: '李四', score: 90 },
  { _id: ObjectId("652d2a7005726f70625ce5d1"), name: '王五', score: 90 }
]
```

2. At time point C, the scores (score) of the three people in the collection test were updated, and zhao liu's information was added.

```
[direct: mongos] tset> db.test.find()
[
  { _id: ObjectId("652d2a6505726f70625ce5cf"), name: '张三', score: 95 },
  { _id: ObjectId("652d2a6b05726f70625ce5d0"), name: '李四', score: 96 },
  { _id: ObjectId("652d2a7005726f70625ce5d1"), name: '王五', score: 97 },
  { _id: ObjectId("652d2beb05726f70625ce5d2"), name: '赵六', score: 90 }
]
```

Detecting Issues and Initiating Key-Based Rollback

At time point D, it was discovered that the scores of Zhang San, Li Si, and Zhao Liu had problems, while Wang Wu's score had no problems. To solve this problem, it was planned to restore the information of Zhang San, Li Si, and Zhao Liu to time point B, which is between time points A and C. Therefore, the _id field of Zhang San, Li Si, and Zhao Liu was used as the filter condition for key-based rollback, and the key-based rollback feature was used to rollback the data to the target time point. Specifically, select _id from the **Specifying Key** list as shown below, enter the following information in the **Input Parameter** box line by line, and initiate a key-based rollback task.

```
ObjectId("652d2a6505726f70625ce5cf")
ObjectId("652d2a6b05726f70625ce5d0")
ObjectId("652d2beb05726f70625ce5d2")
```

Updating Data Steps

Waiting for Key-Based Rollback Task Execution to Complete, Follow These Steps to Update Data

©2013-2025 Tencent Cloud. All rights reserved.

Step 1: Confirming Data in Rollback Collection is Complete and Usable, and Whether There is a Missing Records Collection for the Target Time Point

1. Execute `show collections` to check if a rollback collection or a missing records collection has been generated. As shown below, `test_bak1016151413` is the rollback collection, and `test_bak_missing1016151413` is the missing records collection. This indicates that at the target rollback time point B, there was data for the specified key that had not yet been written (i.e., there was no data for Zhao Liu at this time), otherwise, there would only be a rollback collection.

```
[direct: mongos] tset> show collections
test
test_bak_missing1016151413
test_bak1016151413
```

> **Note:**
> Note: The time at the end of the rollback collection name and the missing records collection name is the operation time, not the target rollback time.

2. Using key-based filter conditions, respectively calculate the number of data in the original test collection, the test_bak1016151413 rollback collection, and the test_bak_missing1016151413 missing collection.

   As shown below, there are 2 in the rollback collection and 1 in the missing records collection. The three pieces of data have consistent `_id` fields with the data in the original collection, and the total number also meets expectations.

```
[direct: mongos] tset> db.test_bak1016151413.find()
[
  { _id: ObjectId("652d2a6505726f70625ce5cf"), name: '张三', score: 90 },
  { _id: ObjectId("652d2a6b05726f70625ce5d0"), name: '李四', score: 90 }
]
[direct: mongos] tset> db.test_bak_missing1016151413.find()
[ { _id: ObjectId("652d2beb05726f70625ce5d2") } ]
```

Step 2: Comparing Data in Rollback Collection and Original Collection to Confirm the Final IDs and Corresponding Keys to Replace Back to the Original Collection, While Filtering Out Unnecessary Records

Check the data in the rollback collection to confirm whether it meets key-based filter conditions and the requirements of returning to target time point B, and select the IDs and corresponding keys of the data to be replaced back to the original collection. As shown below, the 2 pieces of data in the rollback collection meet the requirements of time point B.

```
[direct: mongos] tset> db.test_bak1016151413.find()
[
  { _id: ObjectId("652d2a6505726f70625ce5cf"), name: '张三', score: 90 },
  { _id: ObjectId("652d2a6b05726f70625ce5d0"), name: '李四', score: 90 }
]
```

Step 3: Writing Back Matching Documents to the Original Collection by ID and Corresponding Keys Through Upsert Method Based on the Selected Key List

1. Use the `forEach` command to update the documents corresponding to _id in the rollback collection `test_bak1016151413` to the original collection `test`.

```
db.test_bak1016151413.find({ _id: { $in: [ObjectId("652d2a6505726f70625ce5cf"),
ObjectId("652d2a6b05726f70625ce5d0")] } }).forEach(function (doc) {
  db.test.updateOne({ _id: doc._id }, { $set: doc }, { upsert: true });
});
```

©2013-2025 Tencent Cloud. All rights reserved.

2. Check that the data in the original collection meets expectations. As shown below, the data for Zhang San and Li Si has been updated back to the state of time point B, meeting expectations.

```
[direct: mongos] tset> db.test.find()
[
  { _id: ObjectId("652d2a6505726f70625ce5cf"), name: '张三', score: 90 },
  { _id: ObjectId("652d2a6b05726f70625ce5d0"), name: '李四', score: 90 },
  { _id: ObjectId("652d2a7005726f70625ce5d1"), name: '王五', score: 97 },
  { _id: ObjectId("652d2beb05726f70625ce5d2"), name: '赵六', score: 90 }
]
```

Step 4: Deciding Whether to Delete Corresponding Records in the Original Collection if There is a Missing Records Collection

If there is a missing records collection, it indicates that the target time point does not yet have data for the corresponding IDs in that collection. Depending on the situation, consider whether to delete the corresponding ID data in the original collection. As shown below, Zhao Liu's data does not exist at time point B. If it is determined that this data needs to be cleaned, then perform the data deletion operation.

```
[direct: mongos] tset> db.test_bak_missing1016151413.find()
[ { _id: ObjectId("652d2beb05726f70625ce5d2") } ]
```

After deletion, check that the data in the original collection meets expectations.

```
[direct: mongos] tset> db.test.find()
[
  { _id: ObjectId("652d2a6505726f70625ce5cf"), name: '张三', score: 90 },
  { _id: ObjectId("652d2a6b05726f70625ce5d0"), name: '李四', score: 90 },
  { _id: ObjectId("652d2a7005726f70625ce5d1"), name: '王五', score: 97 }
]
```

©2013–2025 Tencent Cloud. All rights reserved.

Batch Rollback

Last updated：2025-02-08 10:38:31

Overview

Batch rollback refers to rolling back the table data of multiple instances at once to the source instance. Table rollback or flashback to the current instance by key supports batch operations, allowing rollback of tables from multiple instances at once. Newly rolled back tables are named with the suffix _bak. After the rollback is completed, you can modify the database name as needed to improve the efficiency and accuracy of data recovery, avoiding the cumbersome process of manual recovery one by one. Additionally, MongoDB supports viewing all batch rollback history tasks under the current account, helping you understand past operation records in a timely manner and facilitating unified operation and management.

Starting batch rollback task

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the dropdown list, then choose **Batch Rollback**.
3. On the **Rollback Tasks** page, click **Initiate Rollback** to enter the **Batch Rollback Database Table Data** configuration wizard. You can configure batch rollback tasks for database tables and initiate rollback tasks. For detailed operations, see Database Table Rollback.

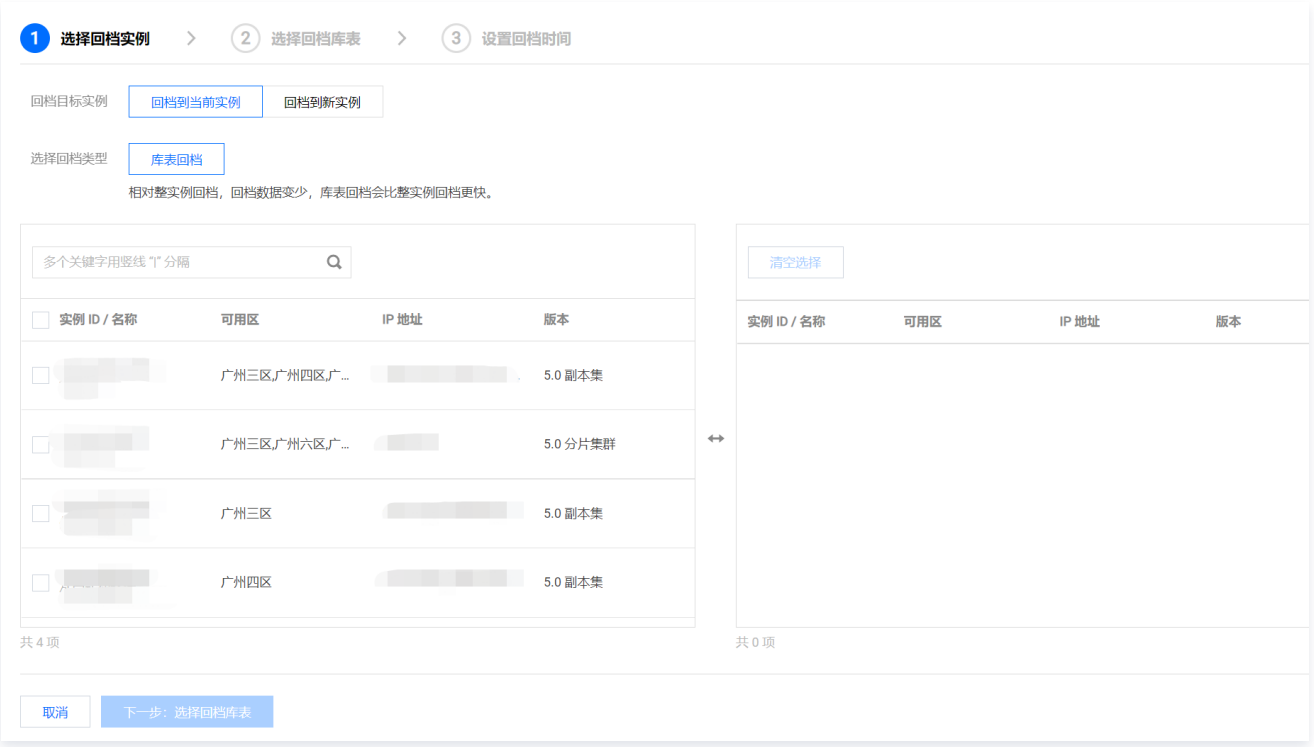

Viewing Batch Rollback Tasks

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the dropdown list, then choose **Batch Rollback**.
3. On the **Rollback Task** page, you can find all batch rollback tasks under the current account, as shown below. You can filter the tasks to view by selecting a time period in the time box.

©2013-2025 Tencent Cloud. All rights reserved.

Bulk Modifying Rollback Database Table Names

Table rollback does not directly rollback data to the original table but creates a backup file. For example, if the source table is test, a new table named test_bak will be created. As shown below, the **rollback table name** is the name of the new table. After the rollback task is completed, you can batch modify the table names as needed.

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the dropdown list, then choose **Batch Rollback**.
3. On the **Batch Rollback Database Table Data** page, under the **Batch Rename Tables** tab, find the rolled back tasks, select instances with table names to be modified one by one, and in the **Operation** column, click **Bulk Modify Table Names**. You will see the information of tables to be modified in the right area, including the original table name, new original table name, rollback table name, and new rollback table name. Confirm the modification by clicking **Batch Rename Tables** at the bottom left to complete the modification, as shown below.

> **Notes:**
> - Batch rename tables, can only modify all tables under a single instance in a single rollback task. If a user initiates a batch rollback task and rolls back the tables of multiple instances, table names must be modified one by one.
> - Batch renaming tables includes modifying the original table name and the rollback table name.
>   - For the original table, add _ori to its original table name.
>   - Rollback table: change the rollback table name to the original table name.

©2013–2025 Tencent Cloud. All rights reserved.

Restoring to Self-built Database

Last updated: 2025-02-08 10:38:52

Restoring a Physical Backup to a Self-built Database

A replica set instance has only one copy of data, while each shard of a sharding cluster has one copy of data. Please restore your data according to your business needs. This document describes how to restore a single copy of data.

Restoring Data to a Single Node

1. Copy the data to an empty data directory in the self-built database, such as /data/27017/.

```
cp -r * /data/27017/
```

2. Restart mongod and check the data. Below is a command sample:

```
./mongod --dbpath /data/27017 --port 27017 --logpath /var/log/mongodb/27017.log --fork
```

Restoring Data to a Replica Set

As a physical backup retains the configuration of the original instance by default, the original configuration needs to be removed first; otherwise, the data may become inaccessible.

1. Restore the data to a single-node self-built database and then restart the node in the form of replica set. Below is a command sample:

```
./mongod --replSet mymongo --dbpath /data/27017 --port 27017 --logpath /var/log/mongodb/27017.log --
fork
```

2. Log in to the node and remove the replica set configuration of the original instance by running the following command:

```
rs.slaveOk()
use local
db.system.replset.remove({})
```

3. Restart the node, add a node to the replica set, initialize it, and check the data. The node added to the replica set should have been started with no data contained. Below is a command sample:

```
rs.initiate({"_id":"mymongo","members":[{"_id":0,"host":"127.0.0.1:27017"},{"_id":1,
"host":"127.0.0.1:27018"},{"_id":2, "host":"127.0.0.1:27019"}]})
```

For the introduction of the rs.initiate() command, please refer to the MongoDB official documentation.

> **Note:**
> Note: Restoring data to a sharded cluster is not supported due to the absence of routing information in the physical backup of the sharded cluster. Therefore, even if the data of each shard is restored to a self-built replica set (each shard of the sharded cluster), mongos can only access the data of the primary shard.

Restoring a Logical Backup to a Self-Built Database

- To ensure the verification after data recovery to a self-built database is not affected, make sure the self-built database is empty.
- For version 3.6, manually remove the config directory and then use the mongorestore command to restore the data of each shard sequentially, as shown below:

©2013-2025 Tencent Cloud. All rights reserved.

```
[root@VM_0_5_centos 1545225029952289395]# ll
total 16
drwxr-xr-x 2 root root 4096 Dec 25 10:38 admin
drwxr-xr-x 2 root root 4096 Dec 25 10:38 config
-rw-r--r-- 1 root root  668 Dec 25 10:38 oplog.bson
drwxr-xr-x 2 root root 4096 Dec 25 10:40 ycsb
[root@VM_0_5_centos 1545225029952289395]# rm -rf config/
[root@VM_0_5_centos 1545225029952289395]# ll
total 12
drwxr-xr-x 2 root root 4096 Dec 25 10:38 admin
-rw-r--r-- 1 root root  668 Dec 25 10:38 oplog.bson
drwxr-xr-x 2 root root 4096 Dec 25 10:40 ycsb
```

- For version 3.2, it is necessary to manually merge the files of a single database table before data can be restored. An example of merging files is as follows:
  In the database `ycsb` directory, there is a table `c_10`, the data files related to this table range from `c_10.bson.gz.chunk-64` to `c_10.bson.gz.chunk-127`, then the merge command is `cat c_10.bson.gz.chunk-* > ./c_10.bson.gz`.

> **Note:**
> In some scenarios of version 3.2, chunk distinction may occur.

Use the mongorestore command to restore data. The -h parameter specifies the self-built database address, the --dir parameter specifies the directory of data files, and the --gzip parameter must be specified to decompress the backup files. The command is as follows:

```
./mongorestore  --gzip --drop -h127.0.0.1:27017 --dir ./1544517027220146694
```

©2013-2025 Tencent Cloud. All rights reserved.

Data Security
Configure Security Group

Last updated: 2025-02-08 10:39:30

You can configure a security group in the TencentDB for MongoDB console to control the outbound/inbound traffic.

Background

Security group is a stateful virtual firewall with filtering feature, used to set the network access control of single or multiple cloud databases. It is an important means of network security isolation provided by Tencent Cloud. The security group is a logical grouping, allowing you to add cloud database instances from the same region with similar network security isolation requirements to the same security group. The cloud database shares the security group list with the CVM and others. The security groups are matched based on rules. See Detailed Description of Security Groups for specific rules and restrictions.

Note

- The CloudDB security group currently only supports network control of VPC private network access and does not support network control of basic networks for the time being.
- Since the CloudDB has no active outbound traffic, the outbound rule does not take effect for the cloud database.
- TencentDB for MongoDB security groups support primary instances, read-only instances, and disaster recovery instances.
- TencentDB for MongoDB supports the security group feature which is implemented based on the allowlist. To use this feature, please submit a ticket.

Directions

Step 1: Create a security group

1. Log in to the Cloud Virtual Machine Console.
2. Select the **Security Group** page in the left sidebar, select a region at the top of the right page and click on **Create**.
3. In the pop-up dialog box, complete the following configurations. Confirm and click on **OK**.

- **Template**: Select the security group template from the dropdown list.
  - **Allow All Ports**: By default, all ports are allowed to the public and private network, which poses certain security risks. Security group rules are added by default. Click on the display template rules below to view the **Outbound Rules** and **Inbound Rules** of this security group template.
  - **Open ports 22, 80, 443, 3389 and ICMP protocol**: Ports 22, 80, 443, and 3389 and the ICMP protocol are opened to the internet by default. All ports are opened to the private network. Security group rules are added by default. The port of TencentDB for MongoDB is 27017 by default. You can ignore this template.
  - **Custom**: After the security group is successfully created, add the security group rules as needed.
- **Name**: Customize the security group name.
- **Project**: By default, select Default Project, but it can be designated to other projects for easier management.
- **Remark**: It is customized and briefly describes the security group for easier management.
- **Advanced Options**: Add a tag to the security group.

4. If **Template** is **Custom**, click on **Set Rules Now** in the **Reminder** dialog box and perform the following steps.

Step 2: Set the security group inbound rules

1. On the **Security Group Rules** page, select the **Inbound Rules** tab and click on **Add Rules**.
2. In the **Add inbound rule** dialog box, set the rules.

- **Type**: Select the default type **Custom**.
- **Source**: Set the source for accessing the database, i.e., inbound source. The following formats are supported to define sources.

Source Format	Format Description

©2013-2025 Tencent Cloud. All rights reserved.

CIDR Notation	The CIDR notation is used for a single IPv4 address or range of IPv4 addresses (such as 203.0.113.0, 203.0.113.0/24, or 0.0.0.0/0, where 0.0.0.0/0 matches all IPv4 addresses). The CIDR notation is used for a single IPv6 address or range of IPv6 addresses (such as FF05::B5, FF05:B5::/60, ::/, or 0::0/0, where ::/ or 0::0/0 matches all IPv6 addresses).
Security Group ID	The security group ID is referenced to match the IP address of the server associated with the security group.
Parameter Template	Reference an IP address object or IP address group object in the Parameter Template.

- **Protocol Port**: Fill in the protocol type and port for the client to access TencentDB for MongoDB. You can view the port information in the Instance List under the **Intranet Address** column, with the default being 27017.
- **Policy**: the default value is Allow.
  - Allow: Access requests of this port are allowed.
  - Reject: Data packets will be discarded without any response.
- **Notes**: It is customized and briefly describes the rule for easier management.

3. Click on **Finish** to complete adding the security group inbound rule.

Step 3: Bind a security group to the instance

> **Note**
> Currently, security groups can be configured only for TencentDB for MongoDB instances in VPC.

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the instance to which you want to bind a security group.
5. In the **Operation** column of the target instance, select **More** > **Security Group**.
   Or click the target instance name, select the **Data Security** tab, and click **Configure Security Group**.
6. In the **Configure Security Groups** dialog box, select the security group to be bound and click on **OK**.

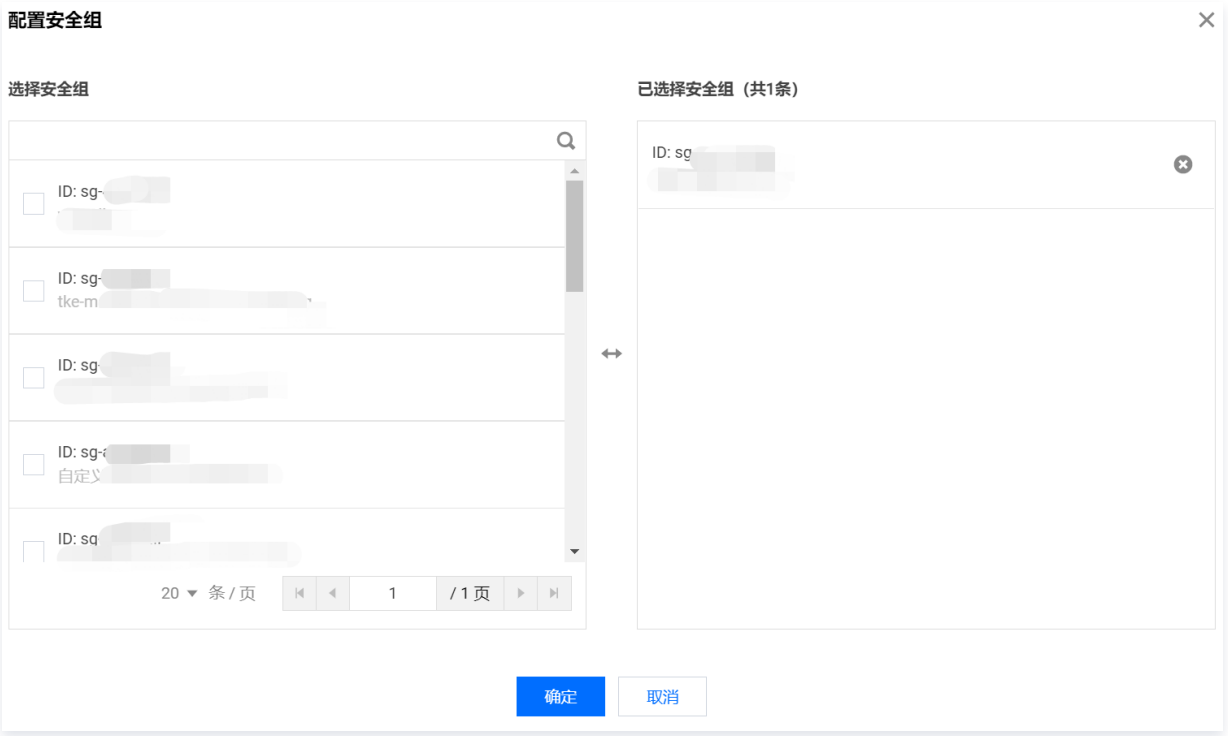

©2013–2025 Tencent Cloud. All rights reserved.

More Operations

Adjusting the Priority of Bound Security Groups

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the instance to which you want to bind a security group.
5. Click the target instance ID, select the **Data Security** tab, and view all current security groups of the instance.
6. Click on **Edit**. You can click on ↑ or ↓ in the **Action** column to adjust the priority of security group filtering.
7. Click on **Save** to complete the addition.

Adjusting the Inbound and Outbound Rules

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the instance to which you want to bind a security group.
5. Click the target instance ID, select the **Data Security** tab, and view all current security groups of the instance.
6. In the security group list, click on **Security Group ID** name to jump to Security Groups page.
7. Find the security group rule to be modified, and click on **Edit** in the **Actions** column to re-edit the security group rule.

Import security group rules

1. On the Security Groups page, select the required security group and click the specific security group ID/name.
2. On the **Inbound Rules** or **Outbound Rules** tab, click on **Import Rules**.
3. In the pop-up dialog box, select the edited inbound/outbound rule template file and click on **Import**.

> **Note**
> - If there are security group rules under the security groups that need to be imported, it is recommended that you export the existing rules first. Otherwise, when importing new rules, the original rules will be overwritten.
> - If there is no security group rule under the security group that needs to import rules, it is recommended that you download the template first and then import the file after editing the template file.

Clone Security Group

1. On the Security Groups page, select **More** > **Clone** in the **Actions** column of the list.
2. In the pop-up dialog box, after selecting the target region and project, click on **OK**.
   If the new security group needs to be associated with CVM, manage CVMs in the security group again.

Deleting a Security Group

1. In the Security Groups page, select the security group to be deleted, and in the Actions column, select **More** > **Delete**.
2. In the pop-up dialog box, click on **OK**.
   If there is a CVM associated with the current security group, you need to remove the security group before deleting it.

More References

For more information about security groups, see Security Groups Overview.

©2013-2025 Tencent Cloud. All rights reserved.

Storage encryption

Last updated: 2025-02-08 10:39:51

Overview

TencentDB for MongoDB provides storage encryption (also known as transparent data encryption, TDE, or static data encryption). It encrypts data before it is written to disk and automatically decrypts it when read into memory. This is implemented by the Database Management Center (DMC) to meet compliance requirements for data encryption.

limitation factor

- The instance type version must be MongoDB 4.4, 5.0, or 6.0.
- The operation account has activated the Key Management Service (KMS). If not activated, you can follow the guided activation process during data encryption setup to activate KMS.
- The operation account must have the `MongoDB_QCSLinkedRoleInKMS` role permission. This role and associated policy grant MongoDB the permission to create and manage keys in KMS. If not authorized, you can follow the guided process during data encryption setup to grant the necessary permissions.

> **Note:**
> - The keys used for encryption are generated and managed by the Key Management Service (KMS). TencentDB for MongoDB does not provide keys or certificates required for encryption.
> - The storage encryption feature incurs no additional charges, but KMS may generate additional fees. For more details, please refer to the Billing Overview.

- If your account has overdue payment, you can't get keys from KMS, which may cause migration, upgrade, and other tasks to fail. For more information, see Payment Overdue.

Notes

- TencentDB for MongoDB uses the AEGIS-256 encryption algorithm for storage encryption.
- After revoking the KMS authorization, a restart will cause the MongoDB database to be unavailable.
- Before enabling storage encryption, automatic backups cannot be physical backups; after enabling, you cannot change automatic backups to physical backups or perform manual physical backups.
- Once enabled, storage encryption cannot be disabled and the key cannot be modified.
- After enabling storage encryption, encrypted data will be inaccessible if the key is disabled or deleted.
- Enabling storage encryption can improve data security but will affect the read/write performance of encrypted databases. Use it based on your actual needs.
- After enabling storage encryption, if your account is in overdue status, you cannot get keys from KMS, which may cause migration, upgrade, and other tasks to fail.
- After enabling storage encryption, existing tables will remain in a non-encrypted state. If you want to encrypt existing tables, it is recommended to create a new encrypted instance and migrate the existing data to the new instance.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The directions for the two types of instances are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, please select Task Management in the left sidebar.
5. In the instance ID/name column of the target instance, click the instance ID to enter the **Instance Details** page.

©2013-2025 Tencent Cloud. All rights reserved.

6. Switch to the **Data Security** page, select the **Storage Encryption** tab, under **Storage Encryption Settings**, in the **Encryption Status** section, click **Click to Activate**.

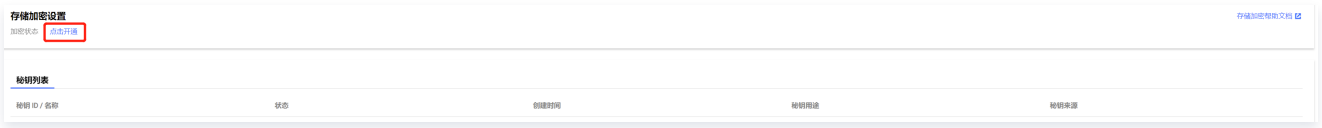

7. In the **Set Data Encryption** window, activate the KMS service and grant KMS key management permissions, and select the key generation method in the **Select Key** option.
   - **Use the key auto-generated by Tencent Cloud**: The key is automatically generated by Tencent Cloud.
   - **Use an existing custom key**: Select a key created on the KMS **Key Management** page, and choose the region of the instance and the name of the key from the dropdown menus below.

> **Note:**
> - To activate storage encryption using a custom key, you need to specify the **Key Usage** as **Symmetric encryption and decryption**. For more information, see Creating a Key.
> - If there are no custom keys, click **Go to create** to create keys in the KMS console. For more information, see Creating a Key.

8. Click **Encryption** to complete the configuration. In the left sidebar, click **Task Management** and wait for the task to complete.

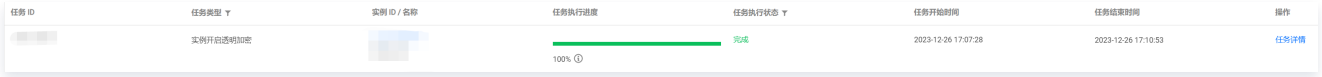

9. Switch to the **Storage Encryption** page, you will see the **Encryption Status** updated to **Activated**, and in the **Key List**, you will see the created keys.

©2013–2025 Tencent Cloud. All rights reserved.

SSL Authentication
Enabling SSL Authentication

Last updated: 2025-02-08 10:40:34

Overview

Secure Sockets Layer (SSL) authentication is a process that authenticates the connection from the user client to the TencentDB server. After SSL encryption is enabled, you can get a CA certificate and upload it to the server. Then, when the client accesses the database, the SSL protocol will be activated to establish an SSL secure channel between the client and the server. This implements encrypted data transfer, prevents data from being intercepted, tampered with, and eavesdropped during transfer, and ultimately ensures the data security for both the client and the server.

Description

Enable SSL, no fees charged, you can try for free.

Note Before Use

- You need to restart the instance to enable SSL. Please perform this operation during off-peak hours, or ensure that your application has a reconnection feature.
- Enabling SSL encryption ensures the security of data access and transfer but will significantly increase CPU utilization. We recommend you enable it only when encryption is required.
- When SSL is enabled, you will receive an expiration alarm 30 days, 15 days, and 7 days before the expiration of your certificate and on its expiration date. Please refresh the SSL certificate in time, or the access authentication through SSL certificate will fail.

Version Description

- New instances: Database versions 4.0 and above support SSL authentication.
- Existing instances: Database version is 3.6, and needs to be upgraded to 4.0 to support SSL authentication.

Prerequisites

- Database instance status: Running, with no other tasks executing.
- The operation is performed in off-peak hours, or the client has an automatic reconnection mechanism.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. In the **Instance ID/Name** column of the target instance, click the instance ID in blue font to enter the **Instance Details** page.
6. Click the **Data Security** tab, then select the **Access Encryption** tab.
7. After **Enable SSL**, click .
8. In the **Enable SSL** window, understand the impact of enabling SSL, and click **Yes**.
9. Wait for the **Enable SSL** status to become **Enabled**, and click **Download Certificate**. If you receive a certificate expiration warning message, and the certificate has expired, please click **Refresh Certificate** to update the certificate file.

©2013-2025 Tencent Cloud. All rights reserved.

10. At the bottom left of the page, obtain the certificate **MongoDB-CA.crt**.
11. To connect to the database via Mongo Shell, see Using Mongo Shell to Connect to Database by SSL Authentication.
    To connect to the database via multi-language SDK, see Using Multi-Language SDKs to Connect to Database by SSL Authentication.

©2013-2025 Tencent Cloud. All rights reserved.

Using Mongo Shell to Connect to Database by SSL Authentication

Last updated: 2025-02-08 10:40:57

Overview

When using Mongo Shell to connect to database, you can enable Secure Sockets Layer (SSL) encryption feature to improve the security of the data linkage. The network connection can be encrypted at the transport layer with the SSL encryption feature to improve the communication data security and ensure data integrity.

Prerequisites

- Create a Linux CVM instance in the same region and VPC as the cloud database MongoDB instance.
- You have obtained the username and password information for database instance access on the **Database Management** page under the **Account Management** tab. For detailed directions, see Account Management.
- You have obtained the private IP and port for database instance access in the **Instance List**. For detailed directions, see Viewing Instance Details.
- SSL encryption has been enabled for the instance. For details, see Enabling SSL Authentication.

Operation Steps

This document uses the Linux operating system as an example to demonstrate the specific operation process.

1. Download the SSL CA certificate. For specific operations, see Enabling SSL Authentication.
2. Upload the certificate file **MongoDB-CA.crt** to the CVM server with Mongo Shell installed.
3. On the CVM server with Mongo Shell installed, execute the following command to connect to the MongoDB database.

> **Note:**
> For MongoDB 4.2 and later, Transport Layer Security (TLS) is used to perform data authentication. TLS is the security protocol of transport layer, an upgraded version of SSL. When you are not sure whether to use SSL authentication or TLS authentication, you can execute `./mongo_ssl -h` to confirm the authentication method.

- **SSL authentication**

```
./bin/mongo -umongouser -plxh***** 172.xx.xx.xx:27017/admin --ssl --sslCAFile MongoDB-CA.crt --
sslAllowInvalidHostnames
```

Please replace the following parameters as needed.

  - -u: refers to the username to connect to the database.
  - -p: refers to the password of the username.
  - 172.xx.xx.xx and 27017 specify the connection IP address (including port number) of the MongoDB instance. If you forget the username and password, please refer to Account Management to view and modify account password information.
  - --sslCAFile: specifies the SSL certificate file path.

- **TLS authentication**:

```
./bin/mongo -umongouser -plxh***** 172.xx.xx.xx:27017/admin --tls --tlsCAFile /data/MongoDB-CA.crt --
tlsAllowInvalidHostnames
```

--tlsCAFile: specifies the TLS certificate file path.

4. The connection success message is shown below.

The prompt messages may vary with different MongoDB shell versions. The following example uses v5.0.15.

©2013-2025 Tencent Cloud. All rights reserved.

More References

For more language SDK connection methods, see Using Multi-Language SDKs to Connect to Database by SSL Authentication.

©2013-2025 Tencent Cloud. All rights reserved.

Using Multi-Language SDKs to Connect to Database by SSL Authentication

Last updated: 2025-02-08 10:41:18

Java

keytool is a native key and certificate management tool in Java, which is convenient for you to manage your public/private keys and certificates for authentication services. Keytool stores keys and certificates in keystore.
Converting certificate format with keytool:

```
keytool -importcert -trustcacerts -file <certificate file> -keystore <trust store> -storepass <password>
```

- `-file <certificate file>` : refers to the SSL certificate or TLS certificate file **MongoDB-CA.crt**.
- `-keystore <trust store>` : specifies the name of the keystore.
- `-storepass <password>` : specifies the password of the keystore.

To set the keystore for JVM system properties, replace trustStore and password with the actual ones to point to the correct keystore. Also, replace the URI concatenation with the user password information to access the database.

```
System.setProperty("javax.net.ssl.trustStore", trustStore);
System.setProperty("javax.net.ssl.trustStorePassword", password);

import com.mongodb.MongoClientURI;
import com.mongodb.MongoClientOptions;

String uri = "mongodb://mongouser:password@10.x.x.1:27017/admin";
MongoClientOptions opt =
MongoClientOptions.builder().sslEnabled(true).sslInvalidHostNameAllowed(true).build();
MongoClient client = new MongoClient(uri, options);
```

Go

The following is a code example of using GO language to connect to the database through SSL authentication. You need to replace the path of the certificate file MongoDB-CA.crt, the account and password, IP information and port information concatenated in the URI as needed.

```
package main

import (
    "context"
    "crypto/tls"
    "crypto/x509"
    "io/ioutil"

     "go.mongodb.org/mongo-driver/mongo"
     "go.mongodb.org/mongo-driver/mongo/options"
)

func main() {
    ca, err := ioutil.ReadFile("MongoDB-CA.crt")
    if err != nil {
        return
    }
    pool := x509.NewCertPool()
    ok := pool.AppendCertsFromPEM([]byte(ca))
    if !ok {
```

©2013-2025 Tencent Cloud. All rights reserved.

```
        return
    }
    tlsConfig := &tls.Config{
        RootCAs:      pool,
        InsecureSkipVerify: true,
    }
    uri := "mongodb://mongouser:password@10.x.x.1:27017/admin?ssl=true"
    clientOpt := options.Client().ApplyURI(uri)
    clientOpt.SetTLSConfig(tlsConfig)

     client, err := mongo.Connect(context.TODO(), clientOpt)
     if err != nil {
         return
     }
    client.Disconnect(context.TODO())
}
```

python

The following is a code example of using Python language to connect database by SSL authentication. You need to replace the path of the certificate file MongoDB-CA.crt, the account and password, IP information and port information concatenated in the URI as needed.

```
from pymongo import MongoClient
uri = "mongodb://mongouser:password@10.x.x.1:27017/admin"
client = MongoClient(uri,
           ssl=True,
           ssl_ca_certs='MongoDB-CA.crt',
           ssl_match_hostname=False)
```

©2013-2025 Tencent Cloud. All rights reserved.

Log Management
Querying Slow Logs

Last updated: 2025-02-08 10:46:47

The TencentDB for MongoDB console supports viewing slow logs. Slow logs record database operation logs with execution times exceeding specified thresholds, used for performance monitoring and optimization, helping developers and operation and maintenance personnel identify and resolve performance bottlenecks.

Background

- In MongoDB, slow logs are often used as a basis to optimize business operations. For more information on slow logs, please refer to the official documentation.
- The system provides two query methods, detailed as follows:
  - Abstract query: Query slow logs by time period, with query results aggregated and analyzed by command type.
  - Specific query: Specify specific operation commands to query slow logs, with query results listing the execution time of operation commands and log details.

Version Description

All versions of MongoDB currently support slow log management.

Must-Knows

- The system logs operations with an execution time of more than 100 ms.
- Slow logs are retained for 7 days, and the time span of a single query does not exceed 1 day.
- The query is limited to the first 10,000 slow logs. If the query is slow, narrow the query time range.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

Querying Slow Log

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the drop-down list, then choose **Replica Set Instance** or **Sharded Instance**. The operations for replica set instances and sharded instances are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the **Target Instance ID** to enter the **Instance Details** page.
6. Select the **Log Management** tab, on the **Querying Slow Logs** page, choose **Query Method** to query slow logs.
   - **Abstract Query**: Select the **Query Time Range**, set the **Time Consumption** threshold, and click **Inquiry**.
   - **Specific Query**: In the **Query Command**, select the specific execution command to be queried, then select the **Query Time Range**, set the **Time Consumption** threshold, and click **Inquiry**.
7. View and analyze slow logs.
   - **Abstract query** results include four fields:
     - **Query Method**: abstract query.
     - **Example sentence**: A statement output with command type as the aggregation dimension, recording operations in the slow log. Users mainly refer to command for troubleshooting.

> **Note:**
> Please pay attention to keywords such as command, COLLSCAN, IXSCAN, keysExamined, and docsExamined. For more log notes, please see MongoDB Official Website.

©2013-2025 Tencent Cloud. All rights reserved.

- command indicates an operation recorded in the slow log.
- COLLSCAN indicates that a full table scan is performed for the query, IXSCAN indicates an index scan.
- keysExamined represents index scan entries, docsExamined represents document scan entries. Larger values of keysExamined and docsExamined indicate no index or low index selectivity. For index optimization, please see Optimizing Indexes to Break Through Read/Write Performance Bottlenecks.

- **Average Execution Time (MS)**: The average execution time of operations aggregated by command type, in milliseconds.
- **Total Counts**: The number of operations aggregated by command type.

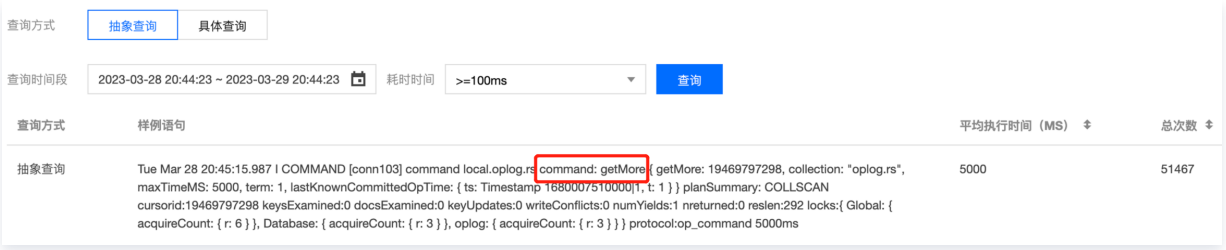

- **Specific Query** results contain three fields:
  - **Query Method**: Specific query.
  - **Time Consumption**: The execution time of business commands, in milliseconds.
  - **Log details**: Business command details.

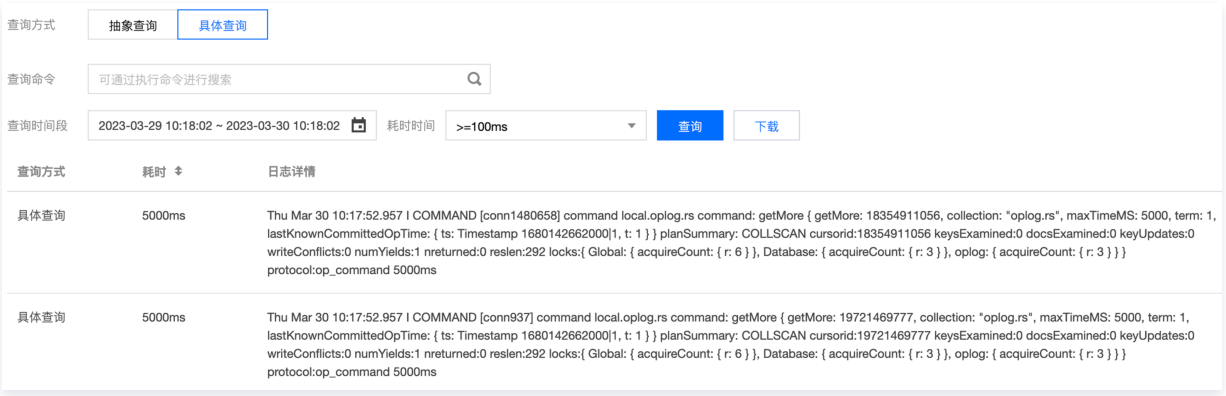

Relevant API

API Interface	API Description
DescribeSlowLogs	Get slow log information
DescribeSlowLogPatterns	Get slow log statistical information

©2013–2025 Tencent Cloud. All rights reserved.

Query Error Log

Last updated: 2025-02-08 10:47:35

The TencentDB for MongoDB console supports querying and managing error logs. Error logs record errors that occur during database operations, including critical errors, general errors, and warnings, for fault troubleshooting and system maintenance, helping to quickly locate and resolve runtime issues.

Must-Knows

The retention time for error logs is 7 days, and the query time span for a single query should not exceed 7 days.

Query Error Log

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the drop-down list, then choose **Replica Set Instance** or **Sharded Instance**. The operations for replica set instances and sharded instances are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the **Target Instance ID** to enter the **Instance Details** page.
6. Select the **Log Management** tab, switch to the **Error Log** page, and you can view the error logs of all instance requests in the current region.

7. At the top of the log list, you can filter the required logs by instance ID, log generation time period, and other keywords.
   - In the drop-down list of **instances**, select the instance ID you want to query.
   - In the time frame area, select the desired time period to query.
   - In the keyword input box, you can filter the required logs by one or more keywords (log category, log details, connection information).

Log Parameters	Log Meaning
Log Type	Supported log types include but are not limited to COMMAND, ACCESS, CONTROL, FTDC, INDEX, NETWORK, QUERY, REPL, SHARDING, STORAGE, RECOVERY, JOURNAL, and WRITE. The specific supported types may vary depending on the MongoDB version. For more information, see Log message.
Log Level	Log levels control the severity of log messages, ranked from most to least serious: FATAL, ERROR, WARNING. • FATAL: Extremely serious errors that usually cause the program or service to stop running. • ERROR: Serious errors that do not immediately cause the program or service to stop running. • WARNING: Warning message indicating potential issues or unexpected behavior that does not immediately affect the normal operation of the system.
Generation Time	Time of error log generation.
Log Details	Describe the original log content. The log content format differs between versions.

©2013-2025 Tencent Cloud. All rights reserved.

	• The log content format before version 4.4 is as follows, including: Timestamp, Log Level, Component, Request Context, Message.
	  ○ The log levels are: I for INFO, W for WARNING, E for ERROR, F for FATAL, D for DEBUG.
	`2014-11-03T18:28:32.450-0500 I NETWORK [initandlisten] waiting for connections on port 27017`
	• The log content format for version 4.4 and later is as follows, including: t (Timestamp), s (Log Level), c (Component), id (Error Code), ctx (Request Context), msg (Message), attr (Additional Information).
	`{"t": {"$date": "2023-03-16T14:33:25.117-04:00"},"s": "I","c": "NETWORK","id": 23016,"ctx": "listener","msg": "Waiting for connections","attr": {"port": 27017,"ssl": "off"}}`

Connection Information	Connection ID information extracted from the request context.
Log ID	Log ID information.

©2013–2025 Tencent Cloud. All rights reserved.

Log Delivery

Last updated：2025-02-08 10:48:25

As customer clusters expand and grow, the reliance on and requirement details for log features continue to rise. The log delivery feature of Cloud Database MongoDB provides an efficient solution: it can automatically capture key slow logs and seamlessly integrate this key data into Cloud Log Service (CLS). Utilizing CLS's powerful analysis and retrieval tools, users can deeply explore performance data, quickly locate and resolve performance bottlenecks, achieving continuous optimization and improvement of database performance.

Supported Versions and Specifications

- Versions: MongoDB 6.0, 4.0
- Architecture Specifications: General Version, Cloud Disk Edition
- Cluster Type: Replica Set, Sharded Cluster

Description

CLS log service is a third-party independent billing cloud product. It supports both **pay-as-you-go (postpaid)** and **resource package (prepaid)** options. You can choose based on your business scenario. For billing standards, see CLS Billing Overview.

Log Shipping Startup Process

To ship logs using CLS, you need to activate the CLS service and authorize the relevant service roles. The startup process is shown in the following diagram.

> **Note:**
> Closing cloud product roles, deleting log topics, and other CLS abnormal statuses may cause shipping exceptions.

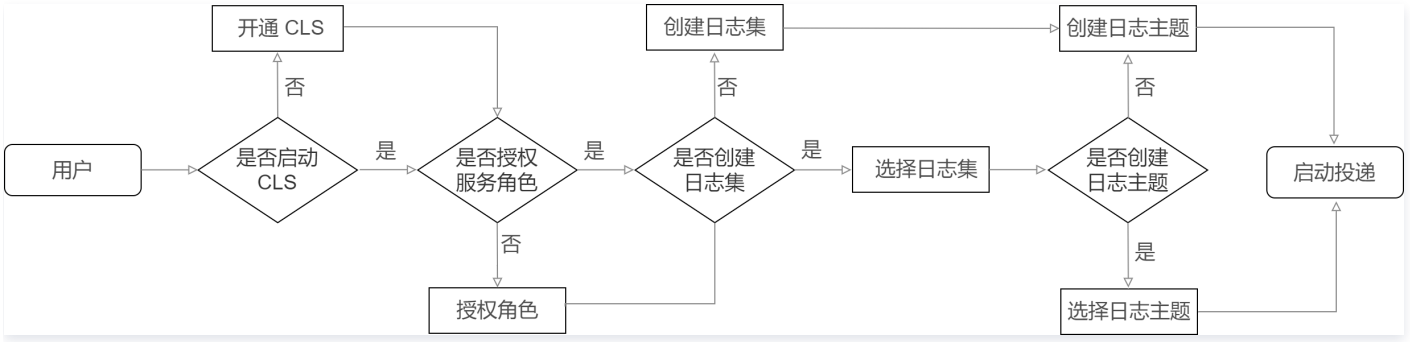

Enabling Operation

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. Select the **Log Management** tab, then select the **CLS Log Shipping** tab.

©2013-2025 Tencent Cloud. All rights reserved.

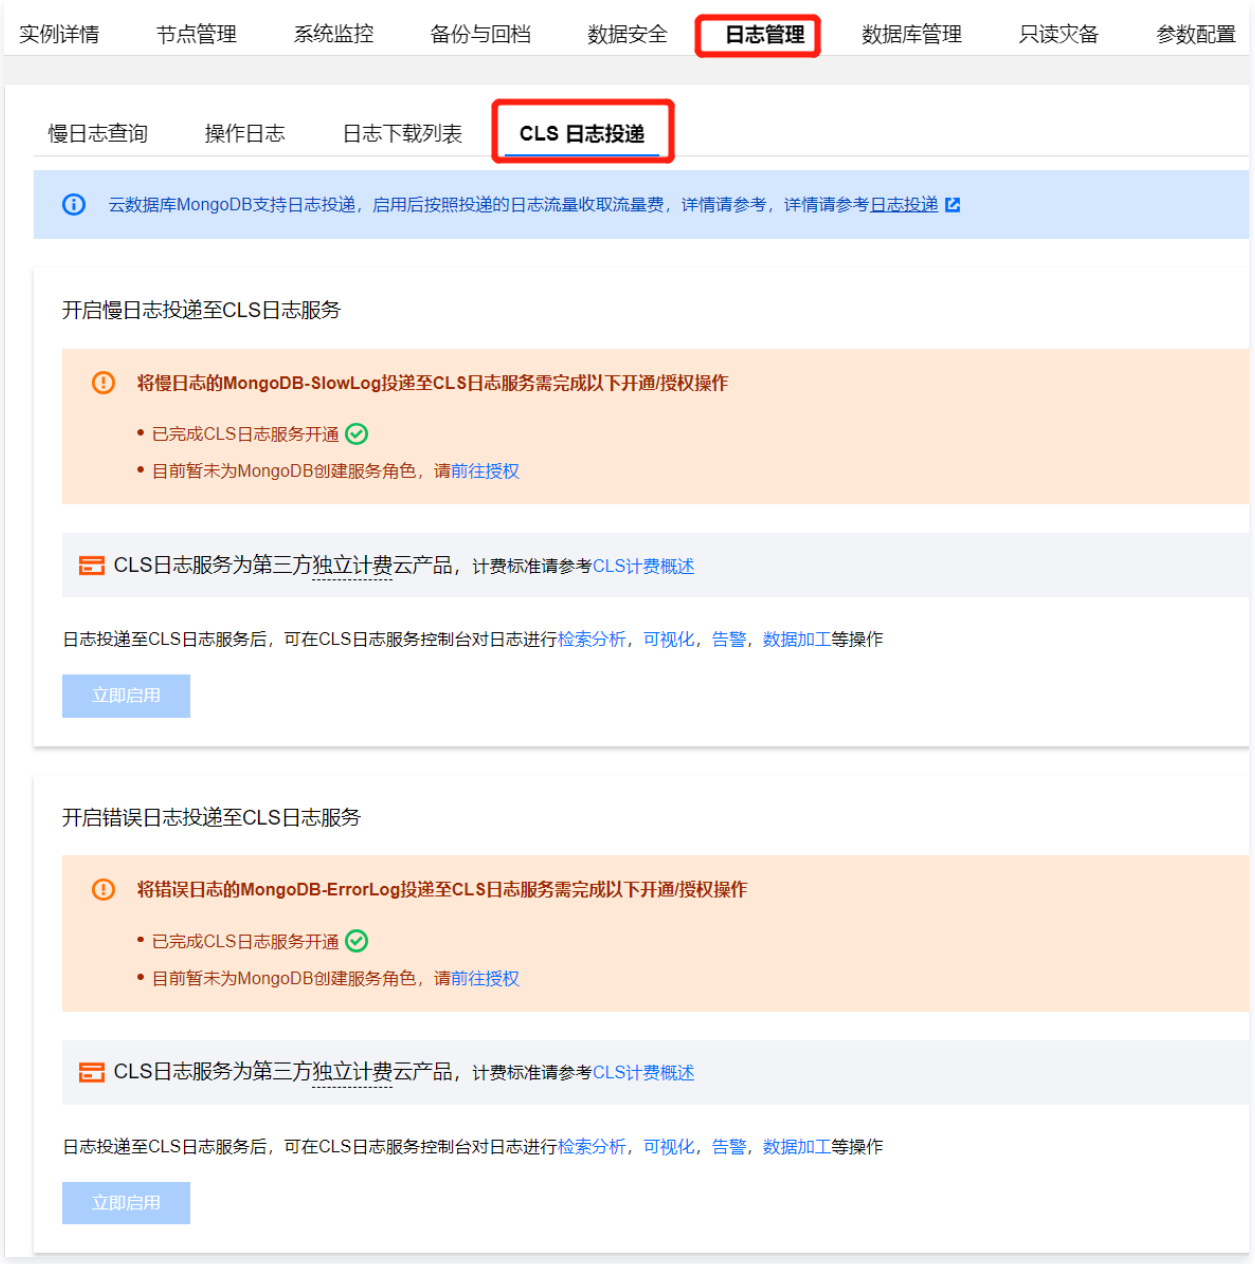

7. TencentDB for MongoDB Log Management will automatically identify whether to enable CLS log service and whether to authorize the service role. Please enable or authorize according to the page prompt information. Slow logs and error logs need to be operated separately.
8. Click **Start Now** to start the log shipping service.
9. In the **Slow Log Delivery** popup window, click **Enable**, configure the log set information for slow log delivery, and click **Yes**.

©2013–2025 Tencent Cloud. All rights reserved.

Parameters	Note:	Example
Region	Select the region for log delivery, support cross-region delivery.	Guangzhou
Logset Operation	A logset is a logical container for log data, used for log classification. A logset can contain multiple log topics. • **Select existing logset**: Select an already created logset. Before enabling slow log delivery, create a logset in the Cloud Log Service (CLS) console. Then, select the existing one. For details, see Creating Logset. • **Creating Logset**: Create a new logset for slow log delivery of TencentDB for Redis.	Creating a logset
Logset	• **Logset Operation** select **Select existing logset**: Please filter the target logset in the dropdown list, perform search, and select the appropriate logset. • **Logset Operation** select **Create logset**: Please edit the logset name in the input box according to the naming requirements to create a new logset. The naming convention is **cloud_*custom name*_logset**, where the custom part only supports English letters, numbers, and underscores, and the length cannot exceed 20.	cloud_redis_slow_logset
Log Topic Operation	Each log topic belongs to a specific logset and is used to distinguish log data from different sources or types within a logset. Log topic is the basic unit for retrieving and analyzing log data. It supports selecting an existing log topic or creating a log topic.	

©2013-2025 Tencent Cloud. All rights reserved.

	• **Select existing log topic**: This can only be set when **Logset Operation** is configured as **Select existing logset**. For log topic related operations, please refer to Log Topic. • **Create log topic**: Create a new log topic under the selected logset.	
Log Topic	• **Log Topic Operation** set to **Select existing log topic**: Please select a log topic under the selected logset from the dropdown list, you can perform search. • **Log topic operation** is set to **Create Log Topic**: Create a new log topic under the selected logset. The naming convention for creating a log topic is **cloud_custom_topic**, where the custom part only supports English letters, numbers, and underscores, and the length cannot exceed 20.	cloud_proxy_slow_log_topic
Retention Date	• Select the retention period for slow log delivery data, which defaults to 30 days. You can choose from 1 to 3600 days. Data will be automatically cleared after the logs expire. • If an existing log topic is selected, the retention period will default to match the retention period of the corresponding existing log topic. • Hover the mouse over ⓘ to see the prompt message. Click **Managing Log Topic** to modify the selected log topic based on the description of the Log Topic.	30
Create an index	Index configuration is a necessary condition for using Cloud Log Service (CLS) for retrieval and analysis. Only when the index is enabled can logs be retrieved and analyzed. It is enabled by default. For specific retrieval and analysis features and methods, refer to Index Configuration. If an existing log topic is selected, the index status will default to the same as the corresponding existing log topic.	Yes

10. After enabling slow log delivery, you can see the slow log delivery status as enabled under log delivery.

©2013–2025 Tencent Cloud. All rights reserved.

Download Log Files

Last updated：2025-02-08 10:49:11

Overview

Download the log files locally for analysis and comparison, which can help users quickly diagnose and solve problems, ensuring stable system operation.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the drop-down list, then choose **Replica Set Instance** or **Sharded Instance**. The operations for replica set instances and sharded instances are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the **Target Instance ID** to enter the **Instance Details** page.
6. Select the **Log Management** tab and download slow logs or error logs as needed.
   - Download slow logs: On the **Slow Query Log** tab, select **Query Method** as **Specific Query**, set the query command and time range, specify the duration, and click **Download** to download the slow logs for the specified time period.

   - Download error logs: On the Error Log tab, click Download to download the error logs for the specified time period.

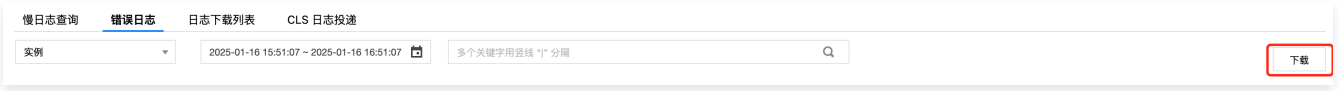

7. Switch to the **Log Download List** tab to view the generated log files as shown below. Click External Network Download to download the log files to your local device for viewing and analysis.

©2013-2025 Tencent Cloud. All rights reserved.

Database Management
Account Management

Last updated: 2025-02-08 10:50:00

The TencentDB for MongoDB console supports creating accounts, setting account permissions, and changing account password information, making it easy to manage database access permissions.

Background

- TencentDB for MongoDB includes two default users: rwuser and mongouser. The 3.2 version system supports rwuser and mongouser users by default, while the 3.6, 4.0, 4.2, and 4.4 version systems default to mongouser users.
  - **rwuser** is the only user that uses MONGODB-CR authentication.
  - **mongouser** and users created in the TencentDB for MongoDB console use SCRAM-SHA-1 authentication.
- Setting multiple accounts and granting each account read and write permissions to different databases allows more granular access to databases and ensures data security.

Version Description

All versions of MongoDB support database account management.

Must-Knows

- Create account and grant access permissions, the system needs 2 minutes for background configuration to take effect.
- It is recommended to reset the database password periodically at an interval of no more than 3 months.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

View Account Information

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. Select the **Database Management** tab to enter the **Account Management** page. You can view all account information of the current database.

Creating Account

1. On the **Account Management** page, click **Create Account**.
2. In the **Account Creation** dialog box, on the **Create Account** tab, configure the account information according to the table below and click **Yes**.

Parameter Name	Required	Parameter Explanation	Parameter Value	Parameter example
Account Name	Yes	Set the name of the new account	New account name setting requirements are as follows: • Character range [1,32]. • Characters in the range [A,Z], [a,z], [1,9], as well as underscores "_" and hyphens "-" can be entered.	test

©2013-2025 Tencent Cloud. All rights reserved.

Account Password	Yes	Set the password for the new account	Password complexity requirements are as follows: • The character range is [8, 32]. • It must contain at least two types of characters from the following: letters, digits, and symbols (!, @, #, %, ^, *, (), _).	test@123
Confirm Password	Yes	Confirm the new account password	Password complexity requirements are as follows: • The character range is [8, 32]. • It must contain at least two types of characters from the following: letters, digits, and symbols (!, @, #, %, ^, *, (), _).	test@123
Remarks	No	Remarks Information	Any character	test
mongouser password	Yes	Enter the mongouser user password	mongouser user's password. Password complexity requirements: • The number of characters is [8, 32]. • Characters can be in the range of [A, Z], [a, z], and [0, 9]. • Allowed special characters include: exclamation mark "!", at "@", hash "#", percent "%", caret "^", asterisk "*", parentheses "()", underscore "_". • Cannot set a single letter or digit.	test@123

3. On the **Set Permissions** page, set the access database permissions for the account.

Parameter Name	Parameter Explanation	Parameter Value
Global Permissions	Set the global permissions for the account to access all databases.	No permission: no read/write data permission. Read-only: read-only data permission. Read/write: read/write data permission.
Instance Details	Set the access specific database permissions for the account.	Inherit global: use global permissions. No permission: no read/write data permission. Read-only: read-only data permission. Read/write: read/write data permission.

4. (Optional) Click **Create New Database**. In the database list, add a new database, enter the new database name in the input box, and click **Yes** to save. Then set the access permissions for the database.

> **Notes:**
> Creating a new database is not a real database, it just presets the database access permissions.

5. Click **Yes** to complete the settings. Wait for 2 minutes for the system configuration to take effect, then you can use the account to access the database.

Modifying Account Permissions

1. In the account list on the **Account Management** tab, find the target account to be modified.
2. In its **Operation** column, click **View/Set**.
3. In the **Set Permissions** dialog box, you can reset the permissions for this account.
4. Click **Yes** to complete the modification.

Modifying Account Password

1. In the account list on the **Account Management** tab, find the target account to be modified.
2. In its **Operation** column, click **Password Modification**.
3. In the **Password Modification** dialog box, reset the **New Password** and **Confirm Password**.
   The password complexity requirements are as follows:
   - Character range [8,32].

©2013-2025 Tencent Cloud. All rights reserved.

- It must contain at least two types of characters from the following: letters, digits, and symbols (!, @, #, %, ^, *, (), _).

4. Click **OK**.

Related Documentation

Viewing the Account URI

1. In the account list on the **Account Management** tab, find the target account to be viewed.
2. In its **Operation** column, click **connection URI**.
3. In the **Connection Help** dialog box, view the URI information of the account connection.
   For more help on connecting instances, refer to Connecting to TencentDB for MongoDB Instance.
4. Click **OK** to close the dialog.

Delete Account

1. In the account list on the **Account Management** tab, find the target account to be deleted.
2. In the **Operation** column, click **Delete**.
3. In the **Delete User** dialog box, confirm the account information to be deleted.
4. Click **OK** to complete the cleanup.

Relevant API

API Name	API Interface
ResetDBInstancePassword	Modify the password of the instance user

©2013–2025 Tencent Cloud. All rights reserved.

Manage slow logs

Last updated: 2025-02-08 10:50:20

The TencentDB for MongoDB console supports viewing slow logs generated during database operation and analyzing slow logs to optimize database performance.

Background

- In MongoDB, slow logs are often used as a basis to optimize business operations. For more information on slow logs, please refer to the official documentation.
- The system provides two query methods, detailed as follows:
  - Abstract query: Query slow logs by time period, with query results aggregated and analyzed by command type.
  - Specific query: Specify specific operation commands to query slow logs, with query results listing the execution time of operation commands and log details.

Version Description

All versions of MongoDB currently support slow log management.

Must-Knows

- The system logs operations with an execution time of more than 100 ms.
- Slow logs are retained for 7 days, and the time span of a single query does not exceed 1 day.
- The query is limited to the first 10,000 slow logs. If the query is slow, narrow the query time range.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Directions

Querying slow log

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** from the drop-down list, then choose **Replica Set Instance** or **Sharded Instance**. The operations for replica set instances and sharded instances are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. Select the **Database Management** tab, then select the **Slow Log Query** tab.
7. On the **Slow Log Query** page, select the **Query Method** to query slow logs.
   - **Abstract Query**: Select the **Query Time Range**, set the **Time Consumption** threshold, and click **Inquiry**.
   - **Specific Query**: In the **Query Command**, select the specific execution command to be queried, then select the **Query Time Range**, set the **Time Consumption** threshold, and click **Inquiry**.
8. View and analyze slow logs.
   - **Abstract query** results include four fields:
     - **Query Method**: abstract query.
     - **Example sentence**: A statement output with command type as the aggregation dimension, recording operations in the slow log. Users mainly refer to command for troubleshooting.

> **Note:**
> Please pay attention to keywords such as command, COLLSCAN, IXSCAN, keysExamined, and docsExamined. For more log notes, please see MongoDB Official Website.
> - command indicates an operation recorded in the slow log.

©2013-2025 Tencent Cloud. All rights reserved.

- COLLSCAN indicates that a full table scan is performed for the query, IXSCAN indicates an index scan.
- keysExamined represents index scan entries, docsExamined represents document scan entries. Larger values of keysExamined and docsExamined indicate no index or low index selectivity. For index optimization, please see Optimizing Indexes to Break Through Read/Write Performance Bottlenecks.

- **Average Execution Time (MS)**: The average execution time of operations aggregated by command type, in milliseconds.
- **Total Counts**: The number of operations aggregated by command type.

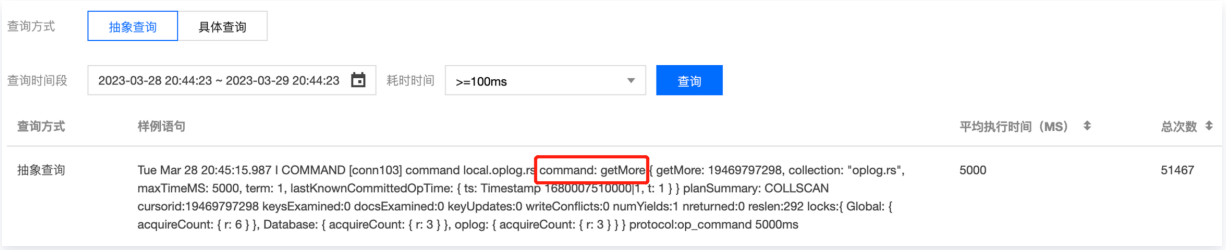

- **Specific Query** results contain three fields:
  - **Query Method**: Specific query.
  - **Time Consumption**: The execution time of business commands, in milliseconds.
  - **Log details**: Business command details.

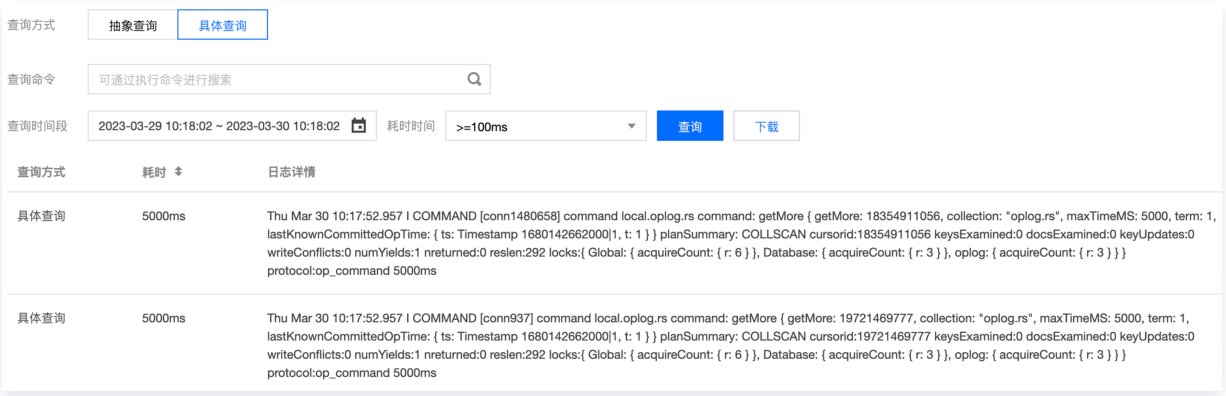

Managing Slow Queries

Viewing slow log request statements

1. On the **Slow Query Management** page, you can view the request statements of the slow logs.
2. In the top right search box, enter query information to search.

Parameter Name	Parameter information
Query statement	Query statement
Op type	Operation Type
Node location	Node where the operation is executed
Command space	Namespace of the database table
Executed time	Time consumption
Details	Details of the execution statement

Batch Killing

1. On the **Slow Log Management** page, select the slow log request statements to be cleared.
2. Click **Batch Kill** at the top of the list to prepare for clearing.
3. In the **Note** dialog box, read the prompt information carefully.

©2013–2025 Tencent Cloud. All rights reserved.

4. Click **OK**.

Downloading Slow Log Files

1. On the **Slow Log Download List** page, you can view the current slow log files.
2. Find the file you need to download, and click **Download** in its **Operation** column.

Relevant API

API Interface	API Description
DescribeSlowLogs	Getting slow log information
DescribeSlowLogPatterns	Getting slow log statistics

©2013–2025 Tencent Cloud. All rights reserved.

Connection Count Management

Last updated: 2025-02-08 10:50:48

Overview

TencentDB for MongoDB will record the Client IP and the number of connections to the current instance. When there are a large number of concurrent application requests, the current database specifications may be unable to meet the demand, and the connection configuration may be insufficient. You can directly increase connections in the console to temporarily solve sudden business expansion needs.

Version Description

- Replica set: All versions of MongoDB support connection management.
- Sharded cluster: MongoDB versions 5.0, 4.4, 4.0, 3.6, and 3.2 support connection management. Version 4.2 does not support it.

Must-Knows

- The system will record the Client IP and the number of connections to the current instance. You can choose to manually release the connection requests.
- If the number of connections reaches or exceeds 80% of the upper limit and affects the establishment of new connections, you can click **Increase Connections** in the console to increase the maximum number of connections to 150% of the original limit for the next 6 hours.
- If the problem persists after increasing the number of connections to 150%, contact after-sales service or submit a ticket for assistance.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Directions

Viewing Connection Usage

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the Instance List, find the target instance.
5. Click the Target Instance ID to enter **Instance Details**page.
6. Select the **Database Management** tab, then select the **Connection Management** tab.
7. View the connection statistics of all clients for the current database.

Parameter Name	Parameter interpretation
Real-time Number of Connections	Statistics Count of All Connection Counts in the Current Database.
Connection Count Proportion	Proportion of Number of Client Connections in the Current Database to Maximum Total Connections.
Maximum Connections	Maximum Limit of Number of Connections.
Current Remaining	Remaining Usage Time to Increase Connection Count.
Client IP	Client IP address of the database connection.
Number of Connections	Connection quantity statistics.

©2013-2025 Tencent Cloud. All rights reserved.

Increasing the Number of Connections

1. On the **Connection Management** page, click **Increase Connections**.
2. In the **Note** dialog box, confirm the prompt information, and click **Yes**.

Relevant API

API Name	API Interface
DescribeClientConnections	Query instance client connection information

©2013–2025 Tencent Cloud. All rights reserved.

Multi-AZ Deployment

Last updated: 2025-02-08 10:51:14

Multi-AZ deployment refers to the deployment of cloud database MongoDB replicas across multiple AZs in the same region. Compared to single-AZ instances (where the primary and replica nodes are in the same AZ), Multi-AZ instances have higher availability and better disaster recovery capability.

Creating Multi-AZ Deployed Instance

1. Log in to the TencentDB for MongoDB purchase page with a Tencent Cloud account.
2. On the purchase page, configure the multi-AZ deployment parameters.

   - In the **Billing Mode** field, select a billing method as needed. Both **Monthly Subscription** and **Pay-as-You-Go** are supported. For more information on how to choose a billing method, see Billing Overview.
   - In the **Region** field, select the region of the multi-AZ deployed instance. We recommend you select the region closest to your end users to reduce latency.
   - In the **AZ** field, click **Multi-AZ Deployment** and select the AZ in the drop-down lists after **Primary Node**, **Secondary Node 1**, and **Secondary Node 2** respectively. To guarantee a smooth cross-AZ switch, multi-AZ deployment does not support deploying most cluster nodes in the same AZ; that is, the primary and secondary nodes can be deployed only in three different AZs separately.
   - For more information on how to configure other parameters, see Creating TencentDB for MongoDB Instance.
3. When selecting **Pay-as-You-Go**, you can click **Billing Details** to view product pricing and confirm the total fees. When selecting **Monthly Subscription**, you can view product pricing and confirm the total fees in Product Pricing.
4. Click **Purchase Now**. After the purchase success message is displayed, click **Go to Console** to enter the instance list page. After the instance status in the **Monitoring/Status** column becomes **Running**, the multiple AZs of the instance will be displayed in the **AZ** column.

Accessing Multi-AZ Deployed Instance

You can use MongoDB Shell or a concatenated URI through the SDK client for multiple programming languages to access a multi-AZ deployed instance. For detailed directions, see Connecting to Instance.

Upgrading from Single-AZ to Multi-AZ Deployed Instance

You can upgrade a single-AZ deployed instance to a multi-AZ deployed instance. For detailed directions, see Modifying Instance AZ.

©2013-2025 Tencent Cloud. All rights reserved.

Disaster Recovery/Read-Only Instances
Overview of Read-Only Disaster Recovery Instances

Last updated: 2025-02-08 10:52:09

Basic Concepts

Read-Only Instances

TencentDB for MongoDB supports creating one or more brand new read-only instances based on the cluster architecture and storage engine of the current instance in the source AZ or other AZs. It automatically synchronizes the data of the current instance to the read-only instances and grants them read-only permission. This helps distribute the read requests of the current instance to the read-only instances, enhancing read/write performance and increasing application throughput.

Disaster recovery instance

TencentDB for MongoDB supports creating one or more brand new disaster recovery instances across regions based on the cluster architecture and storage engine of the current instance. It automatically synchronizes the data of the current instance to the disaster recovery instances and grants them read-only permission. When the region where the current instance is located loses communication due to uncontrollable factors such as power or network issues in any AZ, the disaster recovery instance can be directly promoted to the master instance, enabling cross-region disaster recovery. This quickly supports business needs, helps users enhance business continuity service capability at a lower cost, and ensures data reliability.

Differences Between Read-Only and Disaster Recovery Instances

Both read-only and disaster recovery instances are rebuilt brand new instances based on the cluster architecture and storage engine of the source instance. For specific difference information, please refer to the table below.

Differences	Explanation	Read-Only Instances	Disaster recovery instance
Architecture Type	The system architecture of read-only and disaster recovery instance clusters supports **replica sets** and **sharded clusters**, **single-node** is not supported.	Consistent with the source instance	Consistent with the source instance
Cross-region	Read-only and disaster recovery instances can be built in other regions based on the source instance.	Not supported.	Supported
Cross-AZ	Read-only and disaster recovery instances can be built in other AZs within the same region based on the source instance.	Supported	Supported
Database Version	Mongo-compatible database versions include: 6.0, 5.0, 4.4, 4.2, 4.0.	Consistent with the source instance, cannot be upgraded.	Consistent with the source instance, cannot be upgraded.
Storage Engine	The default storage engine is WiredTiger.	Consistent with the source instance	Consistent with the source instance
Instance Specifications	Requirements for CPU, memory, and disk capacity specifications to ensure the service carrying capacity of read-only and disaster recovery instances.	Cannot be lower than the source instance	Cannot be lower than the source instance

©2013-2025 Tencent Cloud. All rights reserved.

Writing Data	Write data, create or delete database.	Not supported.	Not supported.
Backup and Rollback	Back up data and restore backup data	Not supported.	Not supported.
Account Management	Create or delete database access account	Not supported.	Not supported.
Manually disconnect the association with the source instance	In the console, you can manually disconnect the read-only or disaster recovery instance from the source instance.	Not supported. Only after the original instance is destroyed, the read-only instance will automatically disconnect from the source instance. The read-only instance is promoted to a normal instance and can read and write normally.	Supported. In the console, you can promote the disaster recovery instance, which is converted to a normal instance and can read and write normally, quickly supporting business needs.
Promote AZ	Can the original single-AZ read-only or disaster recovery instance be upgraded to Multi-AZ deployment?	Supported	Supported

How It Works

Read-only and disaster recovery can be understood as a data synchronization service. The synchronization service continuously synchronizes the data and incremental data (databases, tables, indexes, documents, etc.) from the source instance to the target read-only or disaster recovery instance, and restricts the access privileges of the target to read-only, thereby alleviating the pressure on the source cluster.

After creating a read-only disaster recovery instance, full data synchronization is initiated, and the operation log (Oplog) of data synchronization is recorded. Once full data synchronization is completed, incremental synchronization is continuously performed for changed or new data by replaying the Oplog. The general principle is similar to master-slave synchronization within a replica set. The entire process is divided into two phases:

- Full phase, which is the synchronization of full data. Before starting, the latest Oplog timestamp of the current source cluster is recorded. After starting, the metadata, indexes, and data of all databases and tables in the source cluster are read concurrently and inserted into the corresponding databases and tables in the target. The duration of full synchronization is proportional to the data size of the source cluster.
- Incremental phase, which runs after the full phase. It pulls the Oplog of the source cluster based on the Oplog timestamp recorded at the start of the full phase and replays it on the target.

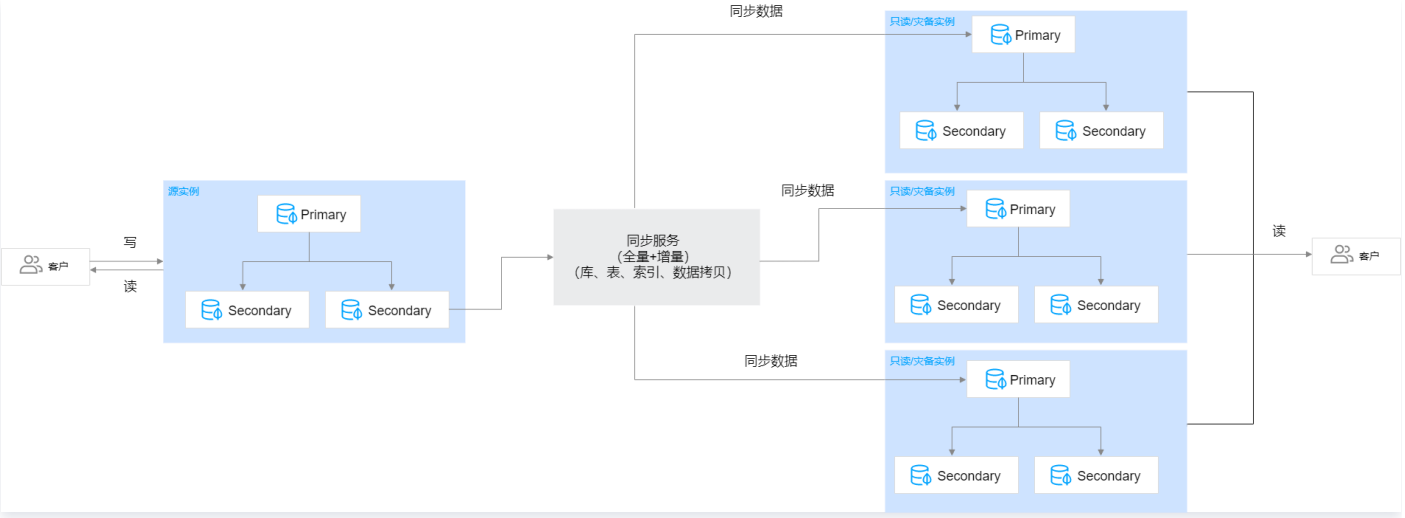

Latency Explanation

©2013-2025 Tencent Cloud. All rights reserved.

Due to data synchronization delay, the real-time synchronization of read-only instances may not be guaranteed. If your business requires read-write separation with high real-time requirements, it is recommended to read from the secondary node of the master instance. For details, see Connecting to TencentDB for MongoDB Instance. Log in to the Console and go to the **Read-only Disaster Recovery** page of the source instance to promptly check the synchronization status and latency from the master instance to the target instance.

Performance Optimization

Similar to the concurrent playback of MongoDB replica set Oplog, the read-only disaster recovery data synchronization service pulls Oplog to the cache, parses Oplog concurrently, hashes by table name to ensure table-level order, then hashes each segment of Oplog by document ID, assigning the Oplog of the same document to the same thread for processing, and performs parallel playback to the target end. This ensures the performance of incremental data synchronization, with synchronization delay reaching the second level.

Data Security

During incremental synchronization, the synchronization service will persist the latest Oplog timestamp that has been synchronized, and the playback process of the synchronization service is idempotent. Therefore, the incremental phase of the synchronization service supports breakpoint resumption. Even if the source or target cluster fails, the synchronization service can ensure data security.
During incremental data synchronization, if the target cluster undergoes a primary-secondary switch due to disk failure, network issues, etc., it may cause data loss. TencentDB for MongoDB has added an Oplog that records synchronization progress. The data synchronization service periodically inserts synchronization progress records into the Oplog stream of the target cluster. After the new primary node takes effect, it will find the latest record in its Oplog stream and resynchronize the data to prevent data loss.

Synchronization Stability

Each read-only disaster recovery instance is supported by a separate data synchronization service. Each data synchronization service uses a distributed lock and lease mechanism to ensure service uniqueness and availability, performs real-time monitoring of synchronization tasks, and scheduled optimization, ensuring stable and reliable data synchronization.

Impacts and Limitations

Impacts on the Source Cluster

The impact of read-only and disaster recovery data synchronization services on the source cluster is limited to the secondary node, using one secondary node (preferably hidden) of the source cluster to retrieve data.

- In the full synchronization phase, getMore requests are used to continuously retrieve data.
- In the incremental synchronization phase, getMore requests are used to continuously retrieve the Oplog.

During both the full and incremental synchronization phases, the synchronization service establishes a sequential read cursor on the secondary node to mark read progress, with minimal impact on the secondary node.

Use Limits

- Read-only and disaster recovery instances belong to the source instance and cannot exist independently.
- Read-only and disaster recovery instances are restricted to be not writable after creation.

  When the source instance is destroyed, the system will automatically disconnect the synchronization service with the source instance. The read-only instance will be upgraded to a normal instance and can read and write normally.
- Database version: For sharded clusters, read-only and disaster recovery instances only support versions 4.0, 4.2, and 4.4; for replica set architecture, versions 3.2, 3.6, 4.0, 4.2, and 4.4 are supported. If the current instance has read-only or disaster

©2013-2025 Tencent Cloud. All rights reserved.

recovery instances, upgrading the database version of the current instance is not allowed. Read-only and disaster recovery instances also cannot upgrade their versions.

- Quantity limit: An instance can create up to 3 read-only and 3 disaster recovery instances.
- Cluster architecture and storage engine: The cluster architecture and storage engine of read-only and disaster recovery instances are fixed to be consistent with the source instance and cannot be modified.
- Account Management: When creating read-only or disaster recovery instances, the account information of the source instance will be automatically synced. Read-only and disaster recovery instances do not support creating or deleting account information. After creating read-only or disaster recovery instances, if the access account and password of the source instance are changed, the account information of the source instance will not be automatically synced. You need to manually modify the account information of the read-only or disaster recovery instances when connecting, otherwise, an error prompt will occur, and the connection will fail.
- Due to network isolation, disaster recovery instances cannot be created between the financial zone and the general disaster recovery domain.
- Backup and Rollback: Read-only and disaster recovery instances do not support backup and rollback data features.
- Data Migration: Read-only and disaster recovery instances do not support migrating data from other instances to read-only or disaster recovery instances.

Sync Limits

- Read-only and disaster recovery synchronization of replica sets is achieved by parsing oplog, and all ddl operations are supported.
- Read-only and disaster recovery synchronization of sharded clusters is achieved by parsing change stream. Change stream wraps an application layer over oplog, providing an API to push data in real-time. Besides basic crud operations, it currently supports ddl operations related to database and table structure indexes, including createIndexes, drop, rename, dropDatabase, create, createIndexes, dropIndexes, collMod, convertToCapped, etc.

Description

The billing method for read-only and disaster recovery instances is the same as the source instance. The billing model can be chosen based on the actual business situation, supporting both prepaid annual and monthly billing and postpaid pay-as-you-go billing. Billing items include computing resources and storage resources. For more information, please refer to Billing Overview.

©2013-2025 Tencent Cloud. All rights reserved.

Create a read-only instance

Last updated: 2025-02-08 10:52:35

Scene Description

In scenarios with more reads but fewer writes, a single instance may not be able to handle the read pressure, which may even affect the main business. Based on the current instance's cluster architecture and storage engine, you can create one or multiple brand new read-only instances in the source instance AZ or other AZs to transfer the current instance's read requests to the read-only instances. This achieves elastic scaling of read capabilities, improves read/write performance, and increases application throughput.

Must-Knows

- Due to the delay in data synchronization, the real-time performance of data synchronization for read-only instances may not be guaranteed. If your business requires read-write separation with high real-time performance, it is recommended to read from the secondary node of the master instance. You can check the synchronization delay between each read-only instance and the master instance in the console.
- The connection method for read-only instances is the same as for master instances. Please refer to Connect Instances.
- During the lifecycle of a read-only instance, it can only read and cannot perform data write update operations.
- Read-only instances do not support manual disconnection from the source instance. They will automatically disconnect from the source instance only when the source instance is destroyed. The read-only instance will then become a standard instance and can perform normal read and write operations.

Version Description

The current MongoDB 4.0, 4.2, 4.4, 5.0, and 6.0 replica set instances all support creating read-only instances. Sharding instances only support version 4.0 and above.

Prerequisites

- The current instance status is running normally, with a large number of read requests, high latency, and slow database performance. For more information, see monitoring indicators.
- The AZ and network of the read-only instance have been planned.
- The storage specifications and purchase quantity of the read-only instance have been estimated.
- The billing mode has been selected based on the business scenario, and the required fees for the read-only instance have been budgeted.

Create a read-only instance

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **NoSQL** > **MongoDB**.
3. In the **MongoDB** drop-down list, select **replica set instance** or **sharding instance**. The operations for replica sets and sharding clusters are similar.
4. Select Region at the top of the Instance List page on the right.
5. In the Instance List, find the target instance.
6. Click the Target Instance ID to enter **Instance Details**page.
7. Select the **Read-Only Disaster Recovery** tab to enter the **Read-Only Instance** page.
8. On the **Read-Only Instance** page, click **Create**.
9. On the **TencentDB for MongoDB Read-Only Instance** purchase page, confirm the **primary instance information** and select the required configuration.

©2013-2025 Tencent Cloud. All rights reserved.

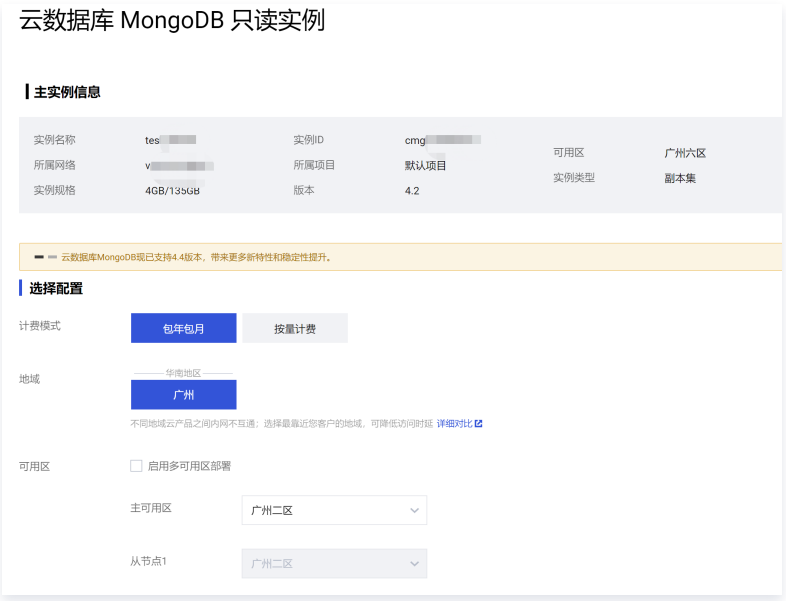

Please refer to the table below to configure the instance specifications according to your actual needs.

Parameter Name	Parameter Description
Billing mode	Both annual and monthly subscription and pay-as-you-go billing are supported. For how to choose a billing method, see Billing Overview.
Region	The region of the read-only instance is fixed to be the same as that of the source instance and cannot be changed.
Availability Zone	Select whether to enable multi-AZ deployment based on actual highly available business requirements.
Database Version	The database version is fixed to be the same as that of the source instance and cannot be changed.
Architecture Type	The architecture type is fixed to be the same as that of the source instance and cannot be changed. For details on the architecture type, see System Architecture.
Storage Engine	The default storage engine is WiredTiger.
Mongod specification	Select the compute specifications of the database instance from the drop-down list. The cpu cores and memory capacity of the read-only instance must be equal to or greater than those of the source instance. Higher specifications result in higher IOPS. For detailed supported specifications, see Product Specifications. After creating the instance, you can adjust the computing specifications of the instance. For specific operations, see Adjusting Instance Configuration.
Mongod shard count	When selecting the sharded cluster architecture type, this parameter is displayed. It is used to set the shard count of the sharded cluster, with a range of [1,20]. The shard count of the read-only instance must be equal to or greater than that of the source instance. Each shard is a replica set. Increasing the shard count can increase the storage capacity of the cluster. Make choice according to your need. After creating the instance, you can adjust the shard count of Mongod. For specific operations, see Adjusting instance configuration.
Disk Capacity	Select the storage capacity of the database instance on the slider. The disk capacity of the read-only instance must be equal to or greater than that of the source instance. The range of disk capacity varies with different Mongod specifications. See Product Specifications. The system defaults to setting the storage space of Oplog to 10% of the selected storage capacity. The size of Oplog can be adjusted in the console instance list. For specific operations, see Adjusting Oplog Capacity. After creating the instance, you can adjust the disk capacity of the instance. For specific operations, see Adjusting instance configuration.

©2013-2025 Tencent Cloud. All rights reserved.

Number of primary/secondary nodes	When the architecture type is a replica set, this parameter is displayed. The default is three nodes (one primary and two secondary), forming a one-primary-two-secondary architecture with three storage nodes. Currently, you cannot customize the number of replicas. You can select five nodes (one primary and four secondary) or seven nodes (one primary and six secondary) from the drop-down list. After creating a read-only instance, you can increase the number of secondary nodes. For specific operations, see Adding Secondary Nodes.
Number of primary/secondary nodes per shard	The architecture type is a sharded cluster, and this parameter is displayed. It is used to set the number of nodes for each shard in the sharded cluster. The system defaults to 3 nodes (1 primary and 2 secondary nodes), meaning each shard has a 3-node architecture of 1 primary and 2 secondary nodes. It supports selecting 5 nodes (1 primary and 4 secondary nodes) or 7 nodes (1 primary and 6 secondary nodes) from the dropdown list, but custom node numbers are not supported. After creating the instance, you can increase the number of secondary nodes per shard. For specific operations, refer to Adding Secondary Node.
Number of Read-Only Nodes	Set the number of read-only nodes, supporting no read-only nodes, 1-5 read-only nodes. Only versions 4.0 and 4.2 support configuring the number of read-only nodes, version 3.6 does not support it. After creating the read-only instance, you can increase the number of read-only nodes. For specific operations, refer to Add Read-Only Node.
Configuration Instructions	Calculate the maximum connections per instance based on the configured Mongod specifications to help you predict whether the current specifications meet your expectations.
Mongos Specs	When the architecture type is selected as a sharded cluster, this parameter is displayed. It is used to configure the specifications of Mongos. After configuring the Mongod specifications, mongos will have default specifications adapted. For example, if Mongod is configured with 2 cores and 4GB, Mongos will be configured with 1 core and 2GB by default. Upgrading mongos specifications will incur charges. For pricing details, refer to Product Pricing. The maximum number of connections for the sharded cluster will be determined by the specifications and number of mongos you choose. You can view the maximum number of connections for the instance in the configuration instructions. After creating an instance, you can adjust the mongos configuration. For specific operations, refer to Adjust the mongos node specification.
Mongos Quantity	When the architecture type is selected as a sharded cluster, this parameter is displayed. It is used to configure the number of mongos. If the instance is deployed in the same AZ, the number of mongos ranges from 3 to 32. If multi-AZ deployment is enabled, the number of mongos ranges from 6 to 32. Increasing the number of mongos will incur charges. For pricing details, refer to Product Pricing. After creating an instance, you can adjust the number of mongos. For specific operations, refer to Adding Mongos Node.
Network Type	Only supports selecting VPC.
IPv4 Network	Select a specific virtual private cloud (VPC) and its subnet. It is recommended that you select the same VPC in the same region as the cloud virtual machine. VPCs have regional attributes (such as Guangzhou), while subnets have availability zone (AZ) attributes (such as Guangzhou Zone 1). A VPC can be divided into one or more subnets. Subnets under the same VPC can communicate with each other by default through the internal network, while VPCs (regardless of whether they are in the same region) are isolated from each other by default on the internal network. After purchasing the instance, you can switch the VPC. For detailed directions, see Switch Network. You can also click **Create VPC** and **Create Subnet** to recreate the required network environment. For detailed directions, see Create a VPC.
IPV6 network	Check whether to enable IPV6 access. Currently, it is not supported.
Security Group	Set security group rules for an instance to control inbound traffic to the database. You can select an existing security group from the dropdown list of selecting an existing security group, or click **Customizing a Security Group** to set new security group inbound rules. For more information, see Configuring a Security Group.
Specified Projects	Allocate a project to an instance. You can manage instances by project.
Tag	Set a tag to an instance. You can manage instances by tag. Click **Add** to select the tag key and tag value.

©2013-2025 Tencent Cloud. All rights reserved.

Instance Name	Set the instance name, default is 500. Please set a recognizable name. Supports Chinese characters, English letters, or digits with a length less than 60, en dash "-", and underscore "_" only.
Quantity	A maximum of 3 read-only instances can be created for an instance.
Purchase period	When selecting the monthly subscription billing mode, you need to choose the purchase duration. The longer the duration, the greater the discount. You can choose based on actual business needs.
Automatic Renewal	When selecting the annual or monthly subscription billing mode, you can choose whether to enable the automatic renewal feature. This means that after the fee expiration, the fee will be automatically deducted from your Tencent Cloud account on a monthly basis. If not enabled, please pay attention to reminder messages and renew timely upon fee expiration. For specific operations, see Renewal Policy.
Total Fees	• When selecting the monthly subscription billing mode, the total cost of the purchased specifications is displayed. • Choose pay-as-you-go to display the hourly fees. Click **Billing Details** to refer to Product Pricing.

10. Confirm that the parameter configuration is correct. Click **Purchase Now**, a success message will appear. Click **Go to Console**, in the instance list, wait for the instance status to display as **Running**. Now you can use it normally.

Viewing Read-Only Instance

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the instance list, find the source instance of the read-only instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.
5. In the **Instance ID/Name** column of the source instance, click the instance ID to enter the **Instance Details** page.
6. Click the **Disaster Recovery/Read-Only Instances** tab and select the **Read-Only Instances** tab.

7. View all read-only instances under the source instance.

Parameters	Parameter Description
Instance ID	Read-only instance ID and its name. Click the instance ID in blue font to jump to the read-only instance details page. For more information, see Viewing Instance Details.
Status	The current running status of the instance, normally: running.
Specification	Instance specifications, including: memory and disk capacity.
Latency	The synchronization status of the read-only instance based on the source instance, and its latency.
Nodes	The number of primary and secondary nodes of the read-only instance.

©2013-2025 Tencent Cloud. All rights reserved.

Network	The VPC name of the read-only instance.
Private network address	The private network IPv4 address allocated by the VPC. When accessing the database, you need to configure the private network IP address and its port information. For specific operations, see Connecting to Instances.
Region	Region and AZ information to which it belongs.
Expiration time.	In the monthly subscription mode, the specific expiration time of the instance is displayed. It is empty in the pay-as-you-go mode.
Operation	Click Configuration Change to adjust the specifications of the read-only instance. For source instance specification adjustment, be sure to synchronously upgrade the specifications of the read-only instance, otherwise data loss may occur.

Relevant API

API Interface	API Explanation
DescribeDBInstances	Querying TencentDB Instance List
RenameInstance	Modify Instance Name
RenewDBInstances	Renews a TencentDB instance

©2013-2025 Tencent Cloud. All rights reserved.

Create a disaster recovery instance

Last updated: 2025-02-08 10:52:56

Scene Description

For scenarios with strong demands for business continuity service and data reliability or regulatory requirements, you can create one or more brand new disaster recovery instances across regions based on the cluster architecture and storage engine of the current instance. If the region where the current instance is located loses communication due to uncontrollable factors such as power or network issues in any AZ, and the high availability (HA) system fails, the disaster recovery instance can be directly promoted to the master instance for cross-regional disaster recovery, ensuring data continuous service capability in a timely manner.

Must-Knows

- Due to the delay in data synchronization, the real-time performance of data synchronization for disaster recovery instances may not be guaranteed. The synchronization delay between each disaster recovery instance and the master instance can be checked in the console.
- During the lifecycle of a disaster recovery instance, it can only be read and cannot perform data write update operations.
- When the source instance to which the disaster recovery instance belongs is destroyed, or the disaster recovery instance is manually promoted, the disaster recovery instance will be converted to a standard instance, which can read and write normally, quickly supporting business needs.

Version Description

Currently, replica set instances of versions 4.0, 4.2, 4.4, 5.0, and 6.0 all support creating read-only instances. Sharding instances only support versions 4.0 and above.

Prerequisites

- The current instance is running normally.
- Planned region and AZ for the disaster recovery instance, and its network.
- Estimated storage specifications and purchase quantity for the disaster recovery instance.
- Selected billing mode based on the business scenario and budgeted for disaster recovery instance costs.

Create a disaster recovery instance

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **NoSQL** > **MongoDB**.
3. In the **MongoDB** drop-down list, select **replica set instance** or **sharding instance**. The operations for replica sets and sharding clusters are similar.
4. Select Region at the top of the Instance List page on the right.
5. In the instance list, find the target instance for which you need to create a disaster recovery instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.
6. Click the Target Instance ID to enter **Instance Details**page.
7. Select the **Disaster Recovery/Read-Only Instances** tab, then go to the **disaster recovery instance** page, and click **Creation**.
8. On the **TencentDB for MongoDB disaster recovery instance** purchase page, confirm the **primary instance information** and select the required configuration.

©2013-2025 Tencent Cloud. All rights reserved.

See the table below to configure instance specifications based on your actual needs.

Parameter Name	Parameter Description
Billing mode	Both annual and monthly subscription and pay-as-you-go billing are supported. For how to choose a billing method, see Billing Overview.
Region	Select the region to which the disaster recovery instance belongs.
Availability Zone	Select whether to enable multi-AZ deployment based on actual highly available business requirements.
Database Version	The database version is fixed to be the same as that of the source instance and cannot be changed.
Architecture Type	The architecture type is fixed to be the same as that of the source instance and cannot be changed. For details on the architecture type, see System Architecture.
Storage Engine	The default storage engine is WiredTiger.
Mongod specification	Select the compute specifications of the database instance from the drop-down list. The CPU cores and memory capacity of the disaster recovery instance must be equal to or greater than those of the source instance. Higher specifications result in higher IOPS. For detailed supported specifications, see Product Specifications. After creating the instance, you can adjust the computing specifications of the instance. For specific operations, see Adjusting Instance Configuration.
Mongod shard count	When the architecture type is a sharded cluster, this parameter is displayed. It is used to set the shard count of the sharded cluster, with a range of [1,20]. The shard count of the disaster recovery instance must be equal to or greater than that of the source instance. Each shard is a replica set. Increasing the number of shards can improve the storage capacity of the cluster. Please select as needed. After creating the instance, you can adjust the shard count of Mongod. For specific operations, see Adjusting instance configuration.
Disk Capacity	Select the storage capacity of the database instance on the slider. The disk capacity of the disaster recovery instance must be equal to or greater than that of the source instance. The disk capacity range varies with different Mongod specifications. See Product Specifications. The system defaults to setting the storage space for Oplog to 10% of the selected storage capacity. The size of the Oplog can be adjusted in the console instance list. For specific operations, see Adjusting Oplog Capacity. After creating the instance, you can adjust the disk capacity of the instance. For specific operations, see Adjusting instance configuration.
Number of primary/secondary nodes	When the architecture type is a replica set, this parameter is displayed. The default is three nodes (one primary and two secondary), forming a one-primary-two-secondary architecture with three storage nodes. Currently, you cannot customize the number of replicas. You can select five nodes (one primary and four secondary) or seven nodes (one primary and six secondary) from the drop-down list. After creating a disaster recovery instance, you can increase the number of secondary nodes. For specific operations, see Adding Secondary Nodes.

©2013-2025 Tencent Cloud. All rights reserved.

Number of primary/secondary nodes per shard	The architecture type is a sharded cluster, and this parameter is displayed. It is used to set the number of nodes for each shard in the sharded cluster. The system defaults to 3 nodes (one primary and two secondary nodes), meaning each shard has a one-primary-two-secondary architecture. You can select 5 nodes (one primary and four secondary nodes) or 7 nodes (one primary and six secondary nodes) from the dropdown list. Customizing the number of nodes is not supported at this time. After creating a disaster recovery instance, you can increase the number of secondary nodes per shard. For specific operations, see Adding Secondary Nodes.
Number of Read-Only Nodes	Set the number of read-only nodes, supporting no read-only nodes, 1-5 read-only nodes. Only versions 4.0 and 4.2 support configuring the number of read-only nodes, version 3.6 does not support it. After creating a disaster recovery instance, you can increase the number of read-only nodes. For specific operations, see Add Read-Only Node.
Configuration Instructions	Calculate the maximum connections per instance based on the configured Mongod specifications to help you predict whether the current specifications meet your expectations.
Mongos Specs	When the architecture type is selected as a sharded cluster, this parameter is displayed. It is used to configure the specifications of Mongos. After configuring the Mongod specifications, mongos will have default specifications adapted. For example, if Mongod is configured with 2 cores and 4GB, Mongos will be configured with 1 core and 2GB by default. Upgrading mongos specifications will incur charges. For pricing details, refer to Product Pricing. The maximum number of connections for the sharded cluster will be determined by the specifications and number of mongos you choose. You can view the maximum number of connections for the instance in the configuration instructions. After creating an instance, you can adjust the mongos configuration. For specific operations, refer to Adjust the mongos node specification.
Mongos Quantity	When the architecture type is selected as a sharded cluster, this parameter is displayed. It is used to configure the number of mongos. If the instance is deployed in the same AZ, the number of mongos ranges from 3 to 32. If multi-AZ deployment is enabled, the number of mongos ranges from 6 to 32. Increasing the number of mongos will incur charges. For pricing details, refer to Product Pricing. After creating an instance, you can adjust the number of mongos. For specific operations, refer to Adding Mongos Node.
Network Type	Only supports selecting VPC.
IPv4 Network	Select a specific virtual private cloud (VPC) and its subnet. It is recommended that you select the same VPC in the same region as the cloud virtual machine. VPCs have regional attributes (such as Guangzhou), while subnets have availability zone (AZ) attributes (such as Guangzhou Zone 1). A VPC can be divided into one or more subnets. Subnets under the same VPC can communicate with each other by default through the internal network, while VPCs (regardless of whether they are in the same region) are isolated from each other by default on the internal network. After purchasing the instance, you can switch the VPC. For detailed directions, see Switch Network. You can also click **Create VPC** and **Create Subnet** to recreate the required network environment. For detailed directions, see Create a VPC.
IPV6 network	Check whether to enable IPV6 access. Currently, it is not supported.
Security Group	Set security group rules for an instance to control inbound traffic to the database. You can select an existing security group from the dropdown list of selecting an existing security group, or click **Customizing a Security Group** to set new security group inbound rules. For more information, see Configuring a Security Group.
Specified Projects	Allocate a project to an instance. You can manage instances by project.
Tag	Set a tag to an instance. You can manage instances by tag. Click **Add** to select the tag key and tag value.
Instance Name	To set the instance name, please set an easy-to-recognize name. It supports only Chinese characters, letters, or digits, hyphens "-", and underscores "_", with a length of less than 60.
Quantity	An instance can have up to 3 disaster recovery instances.
Purchase period	When selecting the monthly subscription billing mode, you need to choose the purchase duration. The longer the duration, the greater the discount. You can choose based on actual business needs.

©2013-2025 Tencent Cloud. All rights reserved.

Automatic Renewal	When selecting the annual or monthly subscription billing mode, you can choose whether to enable the automatic renewal feature. This means that after the fee expiration, the fee will be automatically deducted from your Tencent Cloud account on a monthly basis. If not enabled, please pay attention to reminder messages and renew timely upon fee expiration. For specific operations, see Renewal Policy.
Total Fees	• When selecting the monthly subscription billing mode, the total cost of the purchased specifications is displayed. • Select pay-as-you-go to display the hourly cost. Click billing details for more information, see Product Pricing.

9. Confirm that the parameter configuration is correct. Click **Purchase Now**, a success message will appear. Click **Go to Console**, in the instance list, wait for the instance status to display as **Running**. Now you can use it normally.

Viewing Disaster Recovery Instances

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the instance list, find the source instance associated with the disaster recovery instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.
5. In the **Instance ID/Name** column of the source instance, click the instance ID to enter the **Instance Details** page.
6. Click the **Disaster Recovery/Read-Only Instances** tab and select the **Disaster Recovery Instances** tab.

7. In the disaster recovery instance list, view all disaster recovery instances under the source instance.

Parameters	Parameter Description
Instance ID	Disaster recovery instance ID and its name. Click the instance ID in blue to go to the disaster recovery instance details page. For more information, see Viewing Instance Details.
Status	The current running status of the instance, normally: running.
Specification	Instance specifications, including: memory and disk capacity.
Latency	The status of the disaster recovery instance is synchronized from the source instance and its latency.
Nodes	The number of primary and secondary nodes in the disaster recovery instance
Network	The VPC name associated with the disaster recovery instance
Private network address	The private network IPv4 address allocated by the private network. When accessing the database, you need to configure the private network IP address and its port information. For specific operations, see Connecting to Instances.
Region	Region and AZ information to which it belongs.

©2013-2025 Tencent Cloud. All rights reserved.

Expiration time.	In the monthly subscription mode, the specific expiration time of the instance is displayed. It is empty in the pay-as-you-go mode.
Operation	Click Configuration Change to adjust the specifications of the disaster recovery instance. When adjusting the specifications of the source instance, be sure to synchronously upgrade the specifications of the disaster recovery instance to avoid data loss.

Upgrading a Disaster Recovery Instance to a Master Instance

1. Log in to the TencentDB for MongoDB console.
2. In the MongoDB drop-down list on the left sidebar, select **Replica Set Instance** or **Sharded Instance**. The operations for Replica Set Instance and Sharded Instance are similar.
3. Select Region at the top of the Instance List page on the right.
4. In the instance list, find the source instance associated with the disaster recovery instance.
   - You can find the target instance by entering the instance ID, instance name, private IP address, or tag key in the search box at the top right corner of the instance list.
   - If the instance is not found in the instance list, select **Recycle Bin** in the left navigation bar to confirm whether the instance has been isolated in the recycle bin due to fee expiration. For specific information, see Recycle Bin.
5. In the **Instance ID/Name** column of the source instance, click the instance ID to enter the **Instance Details** page.
6. Click the **Disaster Recovery/Read-Only Instances** tab and select the **Disaster Recovery Instance** tab.
7. In the disaster recovery instance, find the disaster recovery instance to be promoted.
8. Click **Promote** in the **Operation** column. In the **Disaster Recovery Instance Promotion** dialog box, confirm prompt information, and click **Yes**.
9. The disaster recovery instance will be immediately converted to a regular instance and removed from the disaster recovery instance list.

Relevant API

API Interface	API Explanation
DescribeDBInstances	Querying TencentDB Instance List
RenameInstance	Modify Instance Name
RenewDBInstances	Renews a TencentDB instance

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Configuration
Adjusting Database Parameters

Last updated: 2025-02-08 10:53:38

TencentDB for MongoDB allows you to adjust certain database parameters, so that the database features can better adapt to your business needs.

Background

During daily ops, quickly adjusting some database parameters can optimize the database's query and management performance in a targeted manner and adapt to ever-changing business scenes. At the same time, you can view the parameter modification history at any time to ensure that there is a basis for locating exceptions.

Version Description

Currently, MongoDB 3.2 and above support modifying database parameters. However, there are differences in the modifiable parameters on each version as displayed in the console.

Must-Knows

- Currently, the parameter modification feature only supports parameters that can take effect without restart required after modification. It will support other parameters in future iterations. You can also restart the instance on the MongoDB terminal. The restart will cause a disconnection. Therefore, make business arrangements in advance and proceed with caution.
- When updating the cluster architecture or configuration, such as adjusting specifications, nodes, or shards, upgrading nodes, and migrating nodes, you don't need to configure parameters repeatedly, as the system will automatically sync the parameter configuration data.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The instance is running normally.

Operation Steps

Querying parameter configuration

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Replica Set Instance** or **Sharded Instance**. The directions for the two types of instances are similar.
3. In the instance list on the right, find the target instance.
4. Click the Target Instance ID to enter **Instance Details**page.
5. Select the **Parameter Configuration** tab to view the database parameter configuration.

Modifying parameter configuration

1. In the **Modifiable Parameters** tab, click **Modify Running Value**.
2. In the input box of the **Current Running Parameter Value** column, reset the parameter value to be modified. As shown below.

> **Note:**
> - You can modify multiple parameters at the same time.
> - When modifying parameters, be sure to set them according to the **reference values**.
> - In the **Restart upon Modification Required** column, check if the instance will restart. Restarting will cause connection interruptions. Make business arrangements in advance and operate with caution.

©2013-2025 Tencent Cloud. All rights reserved.

The parameter effective range depends on the instance version and architecture. The current version supports the parameters as shown in the table below.

Parameter Name	Restart upon Modification Required	Default Value	Reference Values	Supported Editions	Supported Instance Type	Scope of Application	Parameter Explanation
operation.profiling.slowOpThresholdMs	No	100	[0–65536]	4.0, 4.2, 4.4	Replica set instance, sharded instance	mongod, mongos	Sets the slow query judgment time in ms.
operationProfiling.mode	No	off	[off | slowOp | all]	4.0, 4.2, 4.4	Replica set instance, sharded instance	mongod	This parameter is used to set the database operation performance analysis mode. By setting different modes, you can record the performance data of database operations for performance optimization and troubleshooting. The parameter has the following optional values: • off: Disable operation performance analysis. • slowOp: Record slow operations, i.e., operations that exceed the threshold execution time. The default threshold is 100 milliseconds. • all: Record the performance data of all operations.

©2013–2025 Tencent Cloud. All rights reserved.

setParameter.cursorTimeoutMillis	No	600000	[1–2147483647]	3.2, 3.6, 4.0, 4.2, 4.4	Replica set instance, sharded instance	3.2 and 3.6: mongod; 4.0, 4.2, 4.4: mongod, mongos	This parameter is used to set the maximum idle time for a cursor, meaning the cursor will be auto-disabled if not used within a certain period, releasing related resources. By default, the value of this parameter is 10 minutes. If you need to extend or shorten the cursor timeout, you can modify the value of this parameter. Note that setting this parameter to 0 disables the cursor timeout mechanism, and the cursor will remain open until the client actively closes it.
setParameter.intenalQueryExecMaxBlockingSortBytes	No	33554432	[33554432–268435456]	4.0, 4.2, 4.4	Replica set instance, sharded instance	mongod, mongos	This parameter controls the maximum memory amount MongoDB can use when performing sorting operations. When MongoDB executes queries that require sorting, it may need to sort the result set in memory. If the size of the result set exceeds the set value, MongoDB will use the disk to perform sorting operations, which may lead to performance degradation. Unit: Byte.
setParameter.maxTransactionLockRequestTimeoutMillis	No	5	[0–60]	4.0, 4.2, 4.4	Replica set instance, sharded instance	mongod	This parameter controls the maximum timeout for MongoDB transactions when waiting to acquire a lock. When a transaction needs to acquire a lock, if the lock is held by other transactions, the transaction will wait for a while and try to acquire the lock. If the waiting time exceeds the value set by this parameter, the transaction will throw a timeout exception. Unit: milliseconds.
setParameter.transactionLifetimeLimitSeconds	No	60	[5–300]	4.0, 4.2, 4.4	Replica set instance, sharded instance	mongod	Set the maximum lifespan of a single transaction, unit: seconds. When a transaction starts, MongoDB assigns a unique transaction ID and records the start time of the transaction. If the transaction does not complete within the set time, MongoDB will automatically rollback the transaction and release related resources.
setParameter.failIndexKeyTooLong	No	true	[true | false]	3.2, 3.6, 4.0	Replica set insta	mongod	Sets whether to limit the length of the index key. If this parameter is set to true, when MongoDB creates an

©2013–2025 Tencent Cloud. All rights reserved.

					nce, shard ed insta nce		index, it will throw an error and refuse to create the index if the index key length exceeds the maximum limit supported by MongoDB.
balance.window	No	NUL L	[00:00 | 23:00]	4.0, 4.2, 4.4	Shar ded Insta nce	mongos	The MongoDB cluster balancing operation is used to distribute data as evenly as possible across all nodes in the cluster to improve performance and availability. This parameter specifies within how long MongoDB performs a balancing operation on the cluster to ensure even data distribution among nodes.
openBalance.window	No	fals e	[true | false]	4.0, 4.2, 4.4	Shar ded Insta nce	mongos	Enables or disables the balance window.

3. Click **Yes** to complete the modification.

Querying parameter modification record

1. On the **Parameter Configuration** tab, click **Modification Log**.
2. View the parameter modification log, including values before and after modification, modification status, and modification time.

©2013–2025 Tencent Cloud. All rights reserved.

Creating a parameter template

Last updated: 2025-02-08 10:53:59

The parameter template of TencentDB for MongoDB is a predefined parameter set that meets specific database configuration requirements. Each database version has a pre-configured default parameter template to accommodate general business requirements. Users can also customize appropriate parameter templates based on specific scenario requirements for more efficient database service application.

> **Note:**
> The current parameter template feature is controlled by an allowlist. If needed, please submit a ticket to apply.

<System Default Parameter Template>

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Parameter Template**.
3. On the parameter template page, select **System Default Template**. For each **database version**, TencentDB for MongoDB provides a corresponding default parameter template.

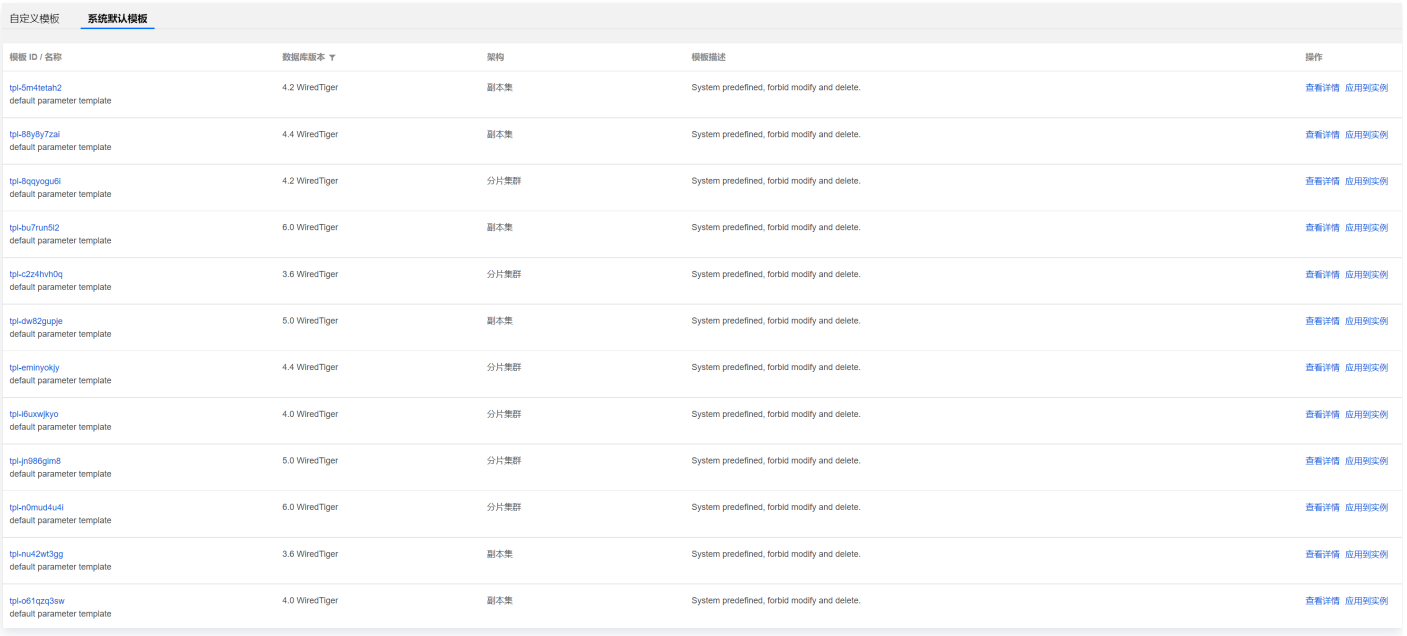

4. Click **Viewing Details** in the **Operation** column to view the specific parameter configuration of the parameter template.

©2013-2025 Tencent Cloud. All rights reserved.

default parameter template 副本集-4.2 WiredTiger

参数名	是否重启生效	参数默认值	参数当前值	参数可修改值
operationProfiling.mode	否		off	[off | slowOp | all]
operation.profiling.slowOpThresholdMs	否		100	[0-65536]
setParameter.cursorTimeoutMillis	否		600000	[1-2147483647]
setParameter.internalQueryExecMaxBlockingSortBytes	否		33554432	[33554432-268435456]
setParameter.maxTransactionLockRequestTimeoutMillis	否		5	[0-60]
setParameter.transactionLifetimeLimitSeconds	否		60	[5-300]

5. Click **Apply to Instances** in the **Operation** column to apply the default parameter template to specific instances. As shown below, select the **Region**, then select one or more specific instances in the region, choose the **Execution Method**, and click **Submit** to apply the default parameter template to the selected instances.

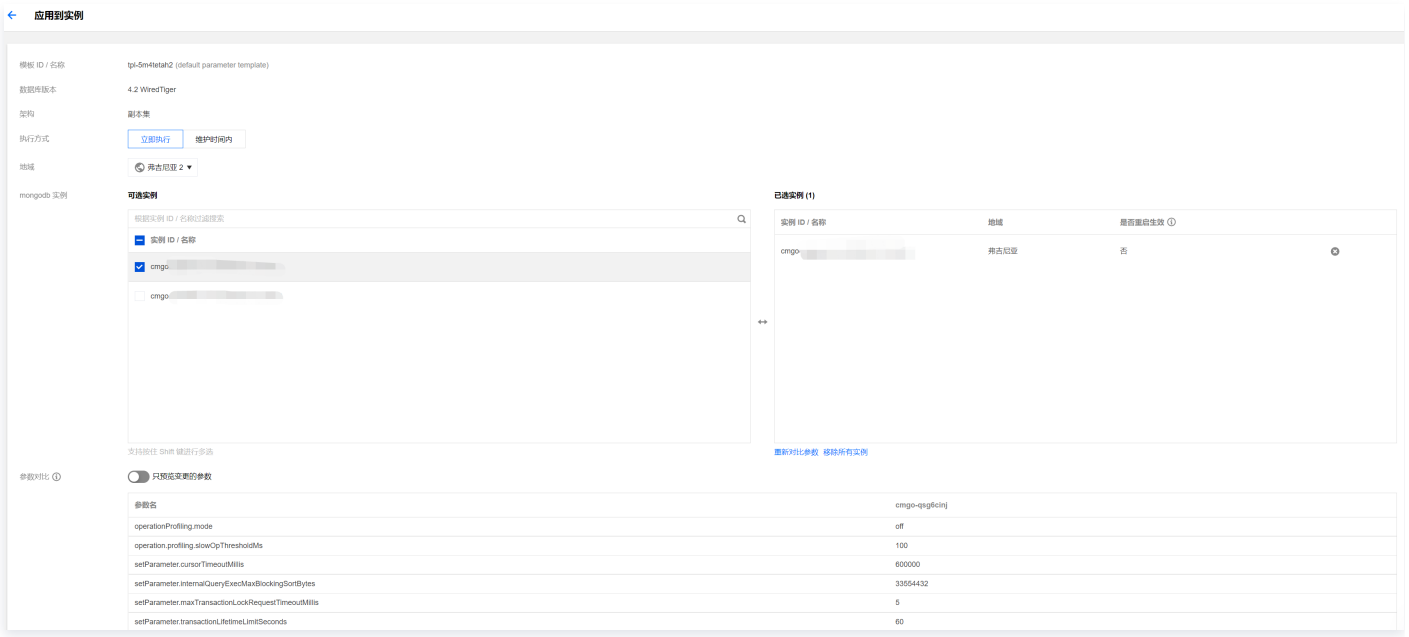

Creating a Custom Parameter Template

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Parameter Template**.
3. On the parameter template page, click create template.
4. In the pop-up **Create Parameter Template** window, configure the following parameters and click **Create and Set Parameters**.
   - **Template Name**: Enter a unique template name.
   - **Database Version**: Select a database version.
   - **Architecture Type**: Select the architecture type for the database instances.
   - **Template Description**: Enter a brief description of the parameter template.

©2013–2025 Tencent Cloud. All rights reserved.

5. Display the parameter list of the template, click **Modifying parameters in batches**, as shown below, in the **Current Value** column, modify the parameter values within the range of **parameter modifiable value**. For parameter meanings, refer to Adjust Database Parameters.

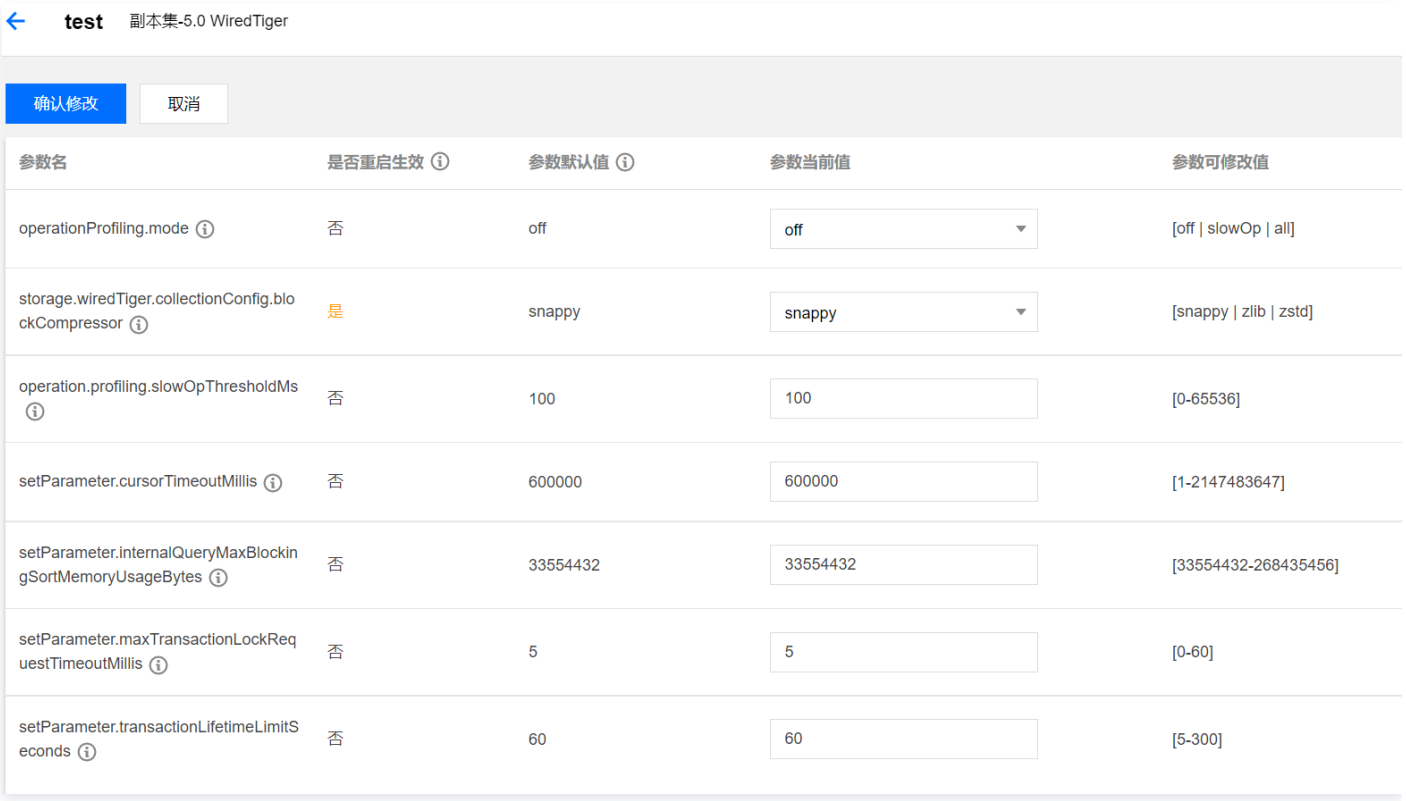

参数名	是否重启生效	参数默认值	参数当前值	参数可修改值
operationProfiling.mode	否	off	off	[off | slowOp | all]
storage.wiredTiger.collectionConfig.blockCompressor	是	snappy	snappy	[snappy | zlib | zstd]
operation.profiling.slowOpThresholdMs	否	100	100	[0-65536]
setParameter.cursorTimeoutMillis	否	600000	600000	[1-2147483647]
setParameter.internalQueryMaxBlockingSortMemoryUsageBytes	否	33554432	33554432	[33554432-268435456]
setParameter.maxTransactionLockRequestTimeoutMillis	否	5	5	[0-60]
setParameter.transactionLifetimeLimitSeconds	否	60	60	[5-300]

6. After modification, click **Confirm Modification**, and the parameter template will be created.
7. (Optional) Return to the parameter list of the template, click **Apply to Instances** to apply the parameter template to specific one or multiple instances.

©2013–2025 Tencent Cloud. All rights reserved.

批量修改参数 | 应用到实例 | 另存为模板

参数名	是否重启生效	参数默认值	参数当前值	参数可修改值
operationProfiling.mode	否	off	off	[off | slowOp | all]
storage.wiredTiger.collectionConfig.blockCompressor	是	snappy	snappy	[snappy | zlib | zstd]
operation.profiling.slowOpThresholdMs	否	100	100	[0-65536]
setParameter.cursorTimeoutMillis	否	600000	600000	[1-2147483647]
setParameter.internalQueryMaxBlockingSortMemoryUsageBytes	否	33554432	33554432	[33554432-268435456]
setParameter.maxTransactionLockRequestTimeoutMillis	否	5	5	[0-60]
setParameter.transactionLifetimeLimitSeconds	否	60	60	[5-300]

8. (Optional) Return to the parameter list of the template, click **Save as Template**, in the **Save as Parameter Template** window, configure the new template name and description, click **Save** to save the current parameter template as a new template.

Applying a Parameter Template to Instances

> **Note:**
> Database instances that have applied a parameter template will not be updated automatically with the template updates. You need to manually reconfigure them by reapplying the updated parameter template to the instances. (Note)

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Parameter Template**.
3. In the right parameter template list, select the template to be applied, and click **Apply to Instance** in the **Operation** column.
4. On the **Apply to Instance** page, select the **Execution Method**, then select the **Region** of the instances, and choose one or more instances from the **Selectable Instances** under **MongoDB instance**.

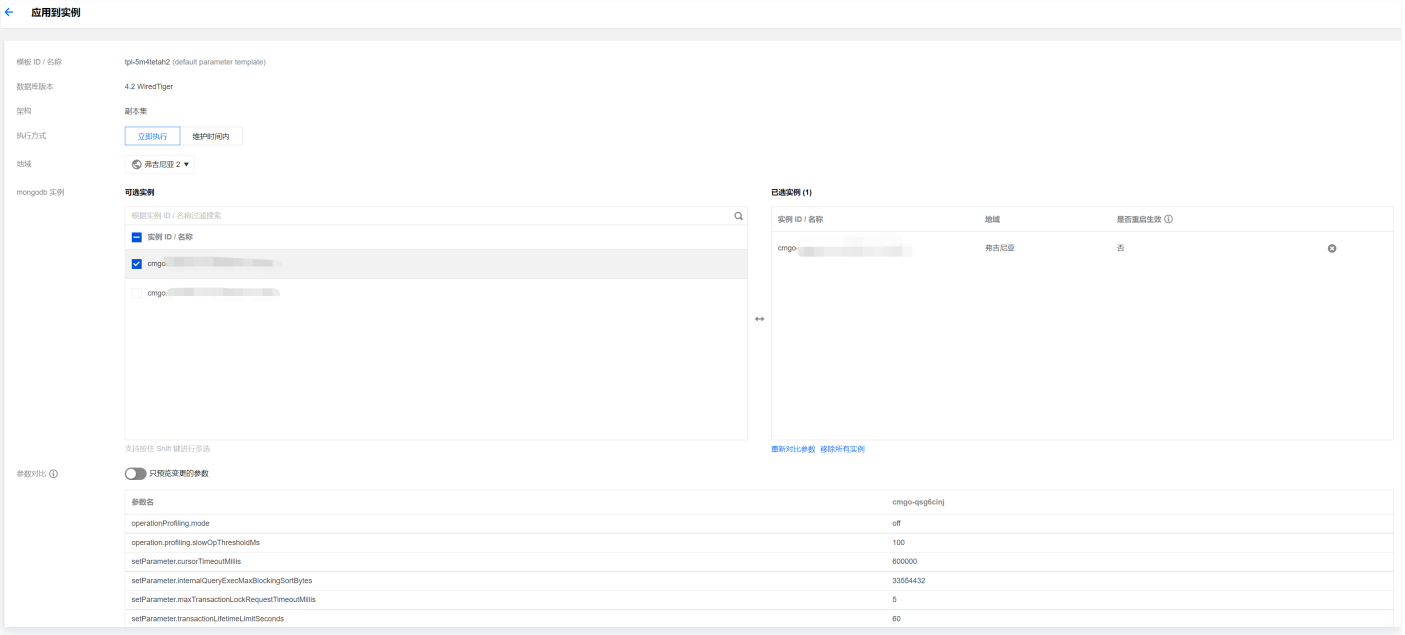

©2013–2025 Tencent Cloud. All rights reserved.

5. Click **Submit**. If the **Execution Method** is set to **Execute Immediately**, the parameter values of the template will be directly applied to the selected instances. If the **Execution Method** is set to **Complete Execution within Maintenance Time**, the task will be executed within the maintenance timeframe, and the parameter values of the template will be applied to the selected instances.

©2013–2025 Tencent Cloud. All rights reserved.

Recycle Bin

Last updated: 2025-02-08 10:54:30

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

Background

<Tencent Cloud recycle bin offers a mechanism for repossessing cloud services. If your account balance is sufficient, after terminating an instance, you can restore it within the retention period if needed.

Version Description

The current MongoDB versions support instance recycling.

Must-Knows

The instance recycling instructions for different billing modes are as follows:

Monthly subscribed instances enter the recycle bin.

- **Retention Duration:** Instances in the recycle bin are retained for 7 natural days.
- **Expiration processing:** If the subscription is not renewed after 7 natural days, the system will release resources and cannot be restored.

> **Note:**
> Seven days before the expiration of cloud service resources, the system will start sending renewal reminder notifications to users. From the eighth day after the expiration, this CloudDB cannot be used anymore and will be recycled in the recycle bin.

Pay-as-you-go instances enter the recycle bin.

- **Retention Duration:** In the case of no arrears, instances actively terminated by users are retained in the recycle bin for 3 days.
- **Expiration processing**: For instances that are not renewed on time, the system will release instance resources after the retention period expires and cannot be restored.

> **Note:**
> - After the account balance is 0, the instance will automatically shut down and stop deduction after 24 hours. The instance will be removed from the instance list and displayed in the recycle bin.
> - For pay-as-you-go instances that have entered the recycle bin, the restoration operation cannot be performed when the account is in arrears. Renew your account first.
> - The pay-as-you-go instance can be saved in the recycle bin for up to three days. Pay attention to the release time and renew and restore them timely.

Prerequisites

- The TencentDB for MongoDB instance has been terminated.
- The balance of the cloud platform account is sufficient.

Operation Steps

You can **renew**, **restore**, and **eliminate** instances in the recycle bin.

©2013-2025 Tencent Cloud. All rights reserved.

Viewing instance in recycle bin

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Recycle Bin**.
3. Select Region at the top of the Instance List page on the right.
4. On the right **Recycle Bin** page, you can see the list of instances in the recycle bin.

Restoring one instance

1. In the recycle bin instance list, find the target instance to be recovered and click **Restore** in the **Operation** column.
2. In the **Restore Instance** window, confirm the instance information to be recovered and click **OK**.
   The instances will return to the instance list in the replica set or sharded cluster from the recycle bin.

Batch restoring instances

1. In the instance list in the recycle bin, select the instances to be restored.
2. Click **Batch Restore** above the list, confirm the instance information in the **Restore Instance** window, and click **OK**.
   The instances will return to the instance list in the replica set or sharded cluster from the recycle bin.

Eliminate instances

1. In the recycle bin instance list, locate the instance you need to shut down and click **Shut down Immediately** in the **Operation** column.
2. In the **Instance Elimination** window, confirm the instance information and click **OK**.

> **Note:**
> After the instance is terminated, all its data will be cleared and cannot be recovered. Back up your data first.

©2013–2025 Tencent Cloud. All rights reserved.

Task Management

Last updated: 2025-02-08 10:54:57

TencentDB for MongoDB allows you to intuitively and quickly track the task execution progress in the console.

Background

Daily OPS involves massive and diverse tasks. Task management can help you quickly and efficiently find tasks and stay up to date with their execution status.

Version Description

Currently, TencentDB for MongoDB 3.2 and above support viewing task execution records.

Prerequisites

- You have applied for a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

You can view the task records and details in the console.

Viewing task record

1. Log in to the TencentDB for MongoDB console.
2. On the left navigation bar, select **MongoDB** > **Task Management**.
3. Select Region at the top of the Instance List page on the right.
4. On the right **Task Management** page, you can see all task records. Hover over the **Task Progress** progress bar to view the task execution process.

Filtering task by time

1. Above the task list, you can select **Today**, **Yesterday**, **Last 7 days**, **Last 30 days**, or choose a time period to filter the tasks to be viewed.
2. In the task list, find the task record to be viewed.

Filtering task by instance name

1. In the search box in the top-right corner of the task list, you can filter the tasks to be viewed by instance name.
2. In the task list, find the task record to be viewed.

Viewing Task Details

1. In the task list, find the task to be viewed and click **Operation** column's **Task Details**.
2. In the **Task Details** window, view the task execution details.
3. After viewing, click **Close**.

©2013-2025 Tencent Cloud. All rights reserved.

Diagnosis and optimization

Last updated: 2025-02-08 10:55:48

TencentDB for MongoDB is connected to the diagnostic optimization feature of DBbrain. The feature monitors and diagnoses database instance exceptions in real time, automatically generates health reports, and gives expert optimization suggestions. This helps you stay on top of the running status of the current database, quickly locate and troubleshoot issues, and promptly optimize the database performance.

Viewing diagnostic optimization

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **Diagnostic Optimization**.
3. At the top of the **DBbrain Diagnosis and Optimization** page, select the target instance in the **Instance ID** drop-down list.

4. View and analyze the diagnosis data of the instance.

Monitoring Type	Monitoring Explanation
Exception Diagnosis	Performs real-time performance monitoring and health inspections on the database and gives diagnosis prompts and optimization suggestions for failures.
Performance Trend	Monitors performance metrics of instances and Mongod nodes, such as resources, requests, and primary/secondary delay.
Real-Time Session	Collects the information of database client sessions in real time, such as the sources and number of sessions as well as the number of active sessions.
Slow Log Analysis	Analyzes the number and duration of slow queries of instances and Mongod nodes in real time.
Space Analysis	Analyzes the database space utilization, including the sizes of data and logs, the daily increase in space utilization, and the estimated number of available days.
MongoStatus	Collects and analyzes the number of requests, updates, deletions, and connections as well as outbound/inbound traffic of the database.
MongoTop	Collects the top data of the database in terms of write operation, read operation, and total request duration.
SQL Throttling	Controls scenarios where excessive CPU resources are consumed due to high traffic. You can create SQL throttling tasks to control the number of access requests and SQL concurrency, thereby ensuring a high service availability.
Index Recommendation	Collects the real-time information of slow queries, automatically analyzes it, and recommends the optimal global index.
Health Report	Scores the instance health based on monitoring metrics and statistics.

©2013-2025 Tencent Cloud. All rights reserved.

Database Auditing Overview

Last updated: 2025-02-08 10:56:23

Overview of Database Audit

Database Audit is a professional, efficient, and comprehensive database audit service independently developed by Tencent Cloud for monitoring database security in real time. It can record the activities of TencentDB instances in real time, manage the compliance of database operations with fine-grained audit, and alert risky database behaviors such as SQL injections and exceptional operations, providing complete and all-around security diagnosis and management features for your TencentDB instances and improving the security of your data assets.
Database Audit can help you deal with the following risks:
Audit Risks Incomplete audit logs make it hard to trace and locate security incidents. Unable to meet the requirements defined by China's Cybersecurity Classified Protection Certification (Level 3). Unable to meet the requirements defined by industry-specific information security compliance documents.
Administrative Risks Misoperation, violation operation, and unauthorized operation by technical personnel harm the safe operation of business systems. Misoperation, malicious operation, and tampering by third-party development and maintenance personnel. Excessive permissions granted to the super admin, which cannot be audited and monitored.
Technical Challenges Database system SQL injections that maliciously pull data from databases and tables. Inability to troubleshoot the sudden increase of database requests that are not slow queries.

Audit Advantages

Comprehensive Audit
Fully records the accesses to databases and executions of SQL statements to meet your audit requirements and ensure database security as much as possible.
Efficient Audit
Different from non-embedded audit mode, Tencent Cloud Database records operations through the database kernel plugin, which makes the records more accurate.
Long-term Retention
Database Audit allows you to retain logs persistently according to your business needs to meet regulatory compliance requirements.
Architecture Characteristics
Database audit adopts the multi-point deployment architecture to guarantee the service availability. It records logs in a streaming manner to prevent tampering and retains them in multiple copies to ensure the data reliability.

Detailed Explanation of Rule-Based Audit

Methods of Audit Rules

- **Full audit**: A full audit of the database's access statements and execution status.
- **Rule-based audit**: Supports setting audit rules for MongoDB database attributes such as **SQL Type**, **Database Name**, **Collection Name**, **Client IP**, **Username**, and audits partial execution statements based on these rules.

Rule-based Audit Operations

- Different types within each rule are additional restriction conditions, i.e., AND (&&) relationship.
- The relationship between rules is OR (||). Each instance can specify one or more audit rules, and as long as any rule is met, it should be audited. For example, Rule A specifies auditing only operations by user1 with execution time >= 1 second, and Rule B audits statements by user1 with execution time < 1 second, so all statements by user1 should be audited.

Database name description

If a statement is of the following table object type:

```
SQLCOM_SELECT, SQLCOM_CREATE_TABLE, SQLCOM_CREATE_INDEX, SQLCOM_ALTER_TABLE,SQLCOM_UPDATE, SQLCOM_INSERT,
```

©2013-2025 Tencent Cloud. All rights reserved.

```
SQLCOM_INSERT_SELECT, SQLCOM_DELETE, SQLCOM_TRUNCATE, SQLCOM_DROP_TABLE
```

For this type of action, the database name is based on the actual operation's database name in the statement. For example, the current database is use db3, and the statement is:

```
select *from db1.test,db2.test;
```

Then db1 and db2 will be used as the target database for rule judgment. If the rule is configured to audit db1, it will be audited. If the rule is configured to audit db3, it will not be audited. If it is not the above table object type statement, the current use database will be used as the target database for judgment. For example, the current database is use db1, and the statement executed is `show databases`, then the current database db1 will be used as the target database for rule judgment. If the rule is configured to audit db1, it will be audited.

Version Description

Currently, TencentDB for MongoDB 4.0, 4.2, 4.4, 5.0, and 6.0 versions support instance auditing.

Description

Database Audit is billed by the amount of audit log storage for every clock-hour, and usage duration shorter than one hour will be calculated as one hour.

Region	Price (CNY/GB/Hour)
China (including finance zones)	0.01
Other countries and regions	0.015

Audit Guidelines

- After audit is enabled for pay-as-you-go cloud databases, when you release a cloud database, its corresponding audit service will also be stopped, and the logs will be automatically deleted and cannot be recovered.
- After audit is enabled for monthly subscribed cloud database instances, when you release an instance or an instance is released upon expiration, its corresponding audit service will also be stopped, and the logs will be automatically deleted and cannot be recovered.

©2013-2025 Tencent Cloud. All rights reserved.

Enabling database audit

Last updated: 2025-02-08 10:56:52

TencentDB for MongoDB provides database audit capabilities, recording database access and SQL statement execution to help enterprises with risk control and improve data security level.

Prerequisites

- You have created a TencentDB for MongoDB instance.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Database Audit**.
3. Above the **Database Audit** page on the right, select the region.
4. In the top right corner of the audit instance list, select an instance whose **Audit Status** is **Disabled**.
5. Click the search box, and in the dropdown list, search for the target instance by **Instance ID**, **Instance Name**, **Tag Key**, or **Tag**.
6. Click the target instance name to enter the **Audit Log** configuration wizard.
7. On the **Enabling Audit Service** configuration wizard tab, view the audit billing description.
8. Select **I agree to the Tencent Cloud Terms of Service** and click **Next**.
9. On the **Audit Service Settings** configuration wizard tab, select the audit **Log Retention Period**, and confirm the required payment after the **Storage Cost**.

> **Note:**
> The audit log retention period can be set to 7 days, 30 days, 3 months, 6 months, 1 year, 3 years, or 5 years. You can also modify it in the console after enabling audit. For more information, see Modifying Log Retention Period. To ensure compliance with security requirements for SQL log retention, it is recommended to select a retention period of 180 days or more.

10. Click **Next**, on the **Audit Rule Settings** configuration wizard tab, choose the rule audit method after **Audit Rules**.
    You can choose **comprehensive audit** or **rule-based audit** as needed, with the default being **comprehensive audit**.
    - **Comprehensive audit**: Audit all statements after enabling.

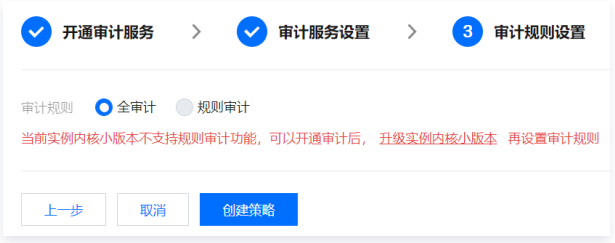

©2013-2025 Tencent Cloud. All rights reserved.

- **Rule-based audit**: Audit database statements based on configured audit items after enabling. Configurable audit items include: sql type, database name, collection name, Client IP, username. Separate multiple items with commas. As shown below.

> **Note:**
> - Multiple database names, collection names, Client IPs, usernames, please use commas to separate them. (Note: Unordered list content)
> - Up to 5 database names, 5 collection names, 5 Client IPs, and 5 usernames can be configured. (Note: Unordered list content)

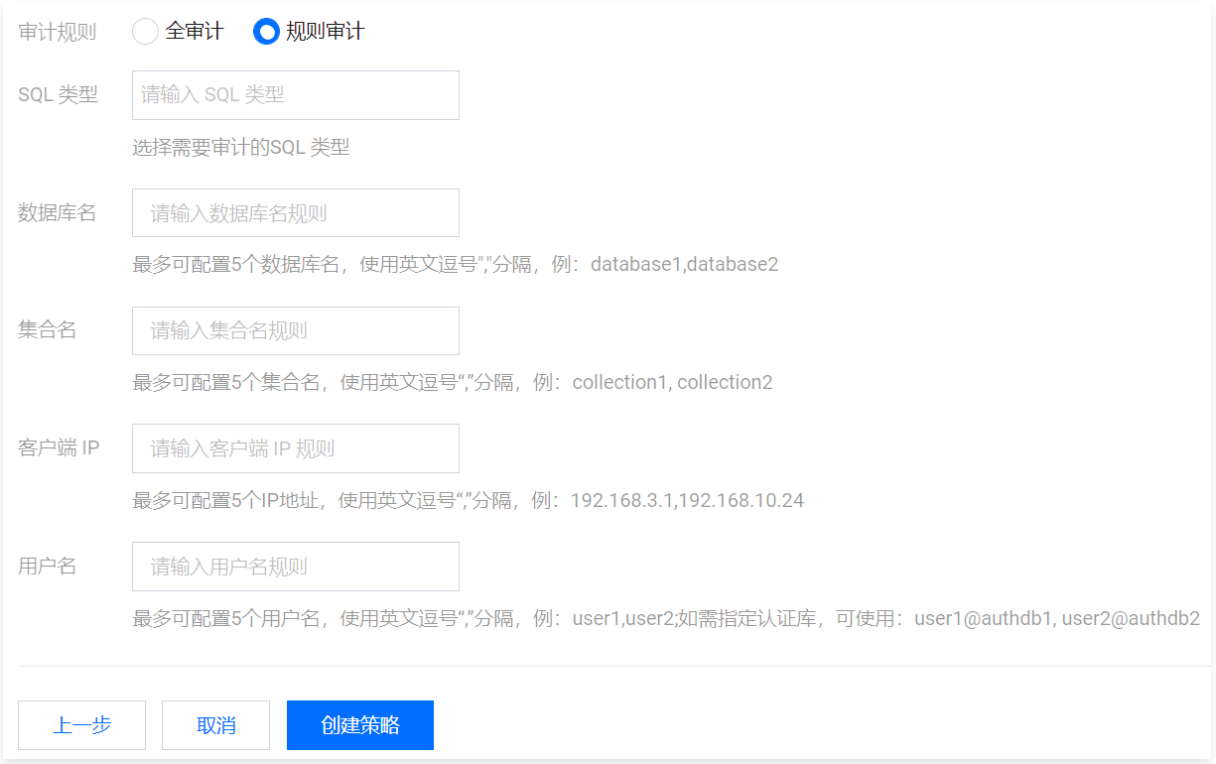

11. Click **Create Policy**, wait for the audit service to be activated and then you can use it. (Ordered list content)

More Operations

- After enabling the audit service, you can analyze the database audit logs at any time to perform risk control. For more information, see Viewing Audit Log. (Unordered list content)
- As business scenarios are constantly changing, it is necessary to adjust audit rules in a timely manner to ensure efficient, accurate, and compliant database supervision. For more information, see Modify Audit Rule. (Unordered list content)
- To manage audit instances, you can view audit instances, close audit instances, modify audit log retention period.

©2013-2025 Tencent Cloud. All rights reserved.

Viewing Audit Instances

Last updated: 2025-02-08 10:57:13

Overview

TencentDB for MongoDB supports displaying instances with activated audit features in the form of a list on the console, to facilitate you based on instance ID and instance name to quickly find audit instances, timely understand the audit rules of audit instances, the storage volume of audit logs, and retention duration.

Prerequisites

- Activated a TencentDB for MongoDB instance with database audit enabled.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Database Audit**.
3. Above the **Database Audit** page on the right, select the region.
4. In the top-right corner of the audit instance list, select an instance whose **Audit Status** is **Enabled**.
5. Click the search box, and in the dropdown list, search for the target instance by **Instance ID**, **Instance Name**, **Tag Key**, or **Tag**.
6. View the audit information of the instance, as shown in the figure below.

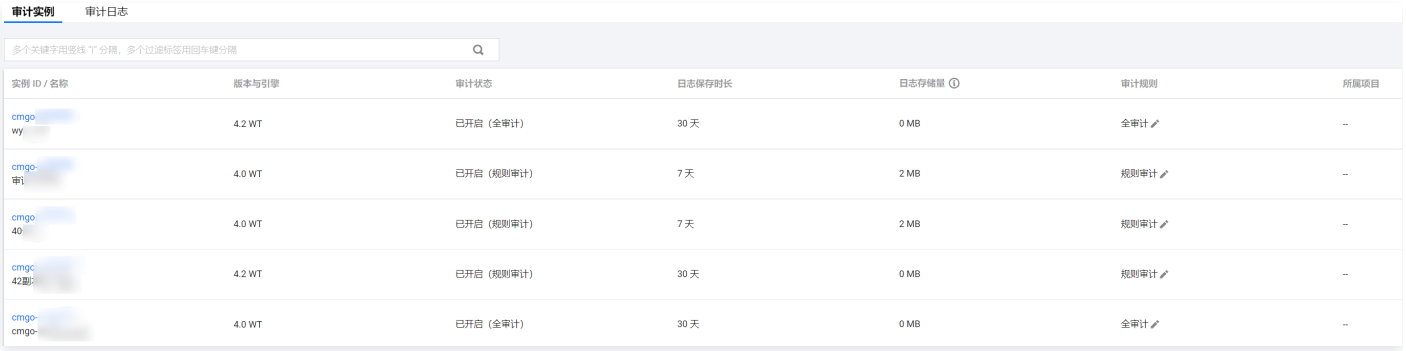

Interface Parameter	<0>Parameter meaning<0>
Instance ID/Name	The ID and name of the audit instance. Click the instance ID to enter the instance details page. For specific information, see Viewing Instance Details.
Version and Engine	Version information and storage engine of the audit instance. WT refers to the WiredTiger engine. For more information, see Versions and Storage Engines.
Audit Status	The audit status of the audit instance, displayed as: Enabled or Disabled.
Log Retention Period	Refers to the duration for which audit logs are retained. To modify, see Modifying Audit Log Retention Period.
Stored Log Size	The storage of audit logs, unit: MB. Data updates may be delayed, and the final data in the billing system shall prevail. This data is for reference only.
Audit Rule	Audit rules of the instance, including comprehensive audit and rule auditing. To modify, please refer to Modify Audit Rule.
Associated project	Project name associated with the instance.
Tag (key: value)	Tag key and tag value associated with the audit instance.

©2013-2025 Tencent Cloud. All rights reserved.

Managing Audit Log

Last updated: 2025-02-08 10:57:35

Overview

After enabling the Database Audit service, the system starts recording operations related to TencentDB for MongoDB. You can view audit log information at any time, including database request access time, Client IP, account name, execution statement, and time consumed by the execution statement.

Must-Knows

- Currently, log file downloads are only available through the Tencent Cloud intranet address. Please download using a Tencent Cloud server in the same region (e.g., download audit logs of database instances in the Beijing region using a CVM in the Beijing region).
- Log files are valid for 24 hours, please download them in time.
- Each database instance can have up to 30 log files. Please download and delete them in time to clean up.
- If the status shows failure, it may be due to too many logs. Please shorten the time window and download in batches.

Prerequisites

- A TencentDB for MongoDB instance has been created, and the instance has enabled Database Audit.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Viewing Audit Logs

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Database Audit**.
3. Above the **Database Audit** page on the right, select the **region**.
4. In the top-right corner of the audit instance list, select an instance whose **Audit Status** is **Enabled**.
5. Click the instance ID with auditing enabled to go to the **Audit Log** page and view the corresponding logs.
6. Select the time period for the audit log in the time frame.
7. In the search box, you can search audit logs based on key information such as **Client IP**, **Account Name**, **Operation Type**, **Execution Time**, **Number of affected rows**, and **Execution Status Code**.

> **Description**
> - **Client IP** address supports using * as a condition for filtering. For example, searching `Client IP: 10.22.23.2*` will match IP addresses starting with 10.22.23.2.
> - **Operation Type**: Selecting the key tag **Operation Type** allows you to perform combination searches. You can check multiple operation types to search simultaneously, and the results will return if any of the conditions separated by | are met.
> - When entering multiple filter tags for search, use the Enter key to separate them. The search results will return audit logs that meet all the filter tags.

8. View audit logs, **Audit Log Fields** are shown in the table below.

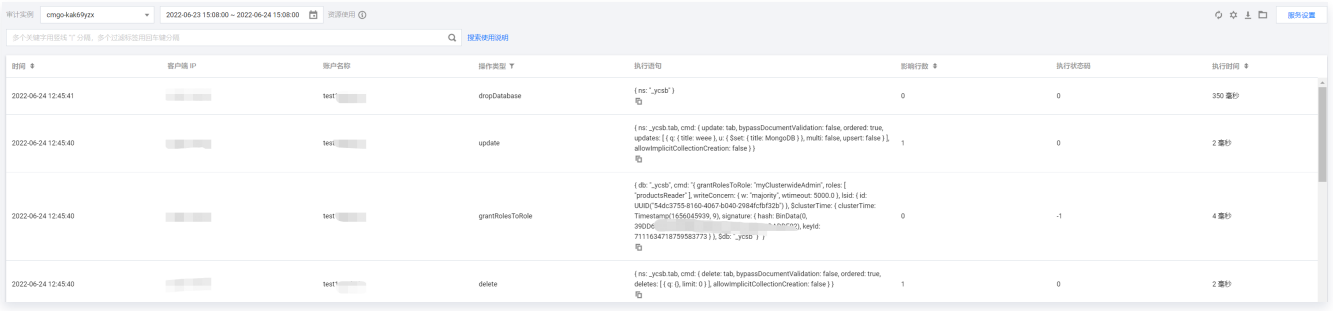

©2013-2025 Tencent Cloud. All rights reserved.

No.	Interface Parameters	Parameter Meaning
1	Time	Operation time of accessing the database.
2	Client IP	Client IP address accessing the database.
3	Account Name	Account name accessing the database.
4	Operation Type	Type of operation statement executed to access the database. Click to select the type to view from the dropdown list.
4	Execution statement.	Request statements executed to access the database.
5	Number of affected rows	Number of changed database rows after executing the operation statement.
6	Execution status code	Returned status code after executing the operation statement. 0 indicates execution success. –1 indicates execution failure. 18 indicates authentication failure. 334 indicates MongoDB authentication protocol unavailable.
7	Execution Time	Time consumed by the operation statement.

Generating and Downloading Audit Log Files

1. In the upper right corner of the audit log list, click .

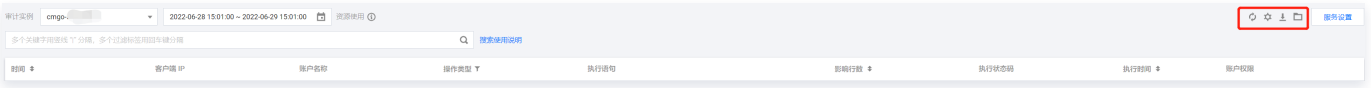

2. In the **Create Log File** dialog box, click **Generate File** to initiate the task of creating a log file.
3. On the **Audit Log File List** page, view the audit log files.

> **Description**
> - Currently, log file downloads are only available through the Tencent Cloud intranet address. Please download through a CVM in the same region (e.g., download the audit log of an instance in the Beijing region through a CVM in the Beijing region).
> - Log files are valid for 24 hours, please download them in time.
> - Each database instance can have up to 30 log files. Please download and delete them in time.
> - If the status shows failure, it may be due to too many logs. Please shorten the time window and download in batches.

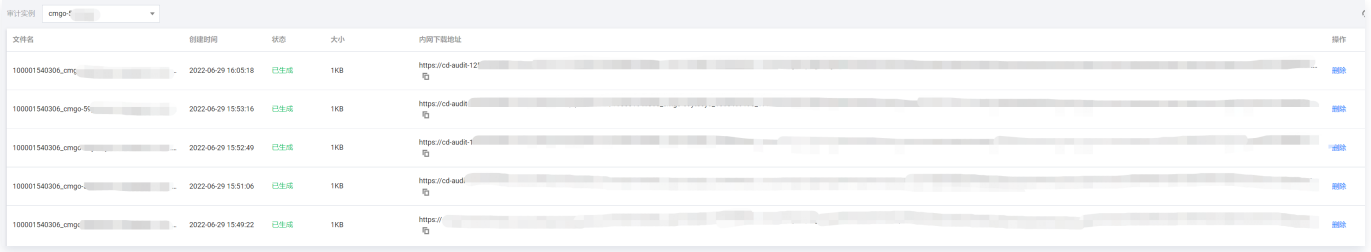

Parameters	Parameter Meaning
File Name	The file name of the audit log is automatically generated by the system.
Creation Time	Time when the audit log file is generated.

©2013–2025 Tencent Cloud. All rights reserved.

Status	Status of the task generating the audit log file, including: generating, generated.
Size	Size of the audit log file.
Private network download address.	Copyable internal network address to download log files.
Operation	Click delete to clear audit log files.

More Operations

Following actions allow you to manage the audit log list.

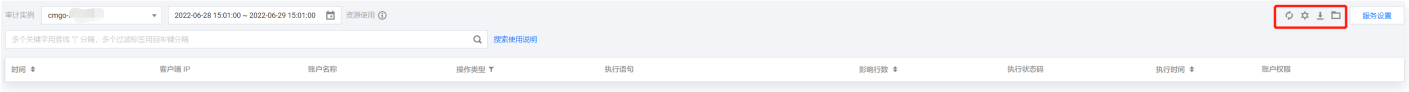

- Click to customize audit log list columns.
- Click to refresh the audit log list.
- Click to directly access the audit log file list page.

©2013–2025 Tencent Cloud. All rights reserved.

Modify Audit Rule

Last updated: 2025-02-08 10:58:14

Overview

TencentDB for MongoDB audit service supports adjusting audit rules at any time according to business scenarios to ensure efficient, accurate, and compliant database supervision.

Prerequisites

- Activated a TencentDB for MongoDB instance with database audit enabled.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Database Audit**.
3. Above the **Database Audit** page on the right, select the region.
4. In the top-right corner of the audit instance list, select an instance whose **Audit Status** is **Enabled**.
5. Click the search box, and in the dropdown list, search for the target instance by **Instance ID**, **Instance Name**, **Tag Key**, or **Tag**.
6. In the **Audit Rule** column of the target instance, click  .
7. In the **Audit Rule** panel displayed on the right, you can modify the **audit rule** method, and reselect **comprehensive audit** or **rule-based audit**.

   If it is **rule-based audit**, please reset the audit rules according to the interface prompt information.

> **Note:**
> - Multiple database names, collection names, Client IPs, usernames, please use commas to separate them. (Note: Unordered list content)
> - Up to 5 database names, 5 collection names, 5 Client IPs, and 5 usernames can be configured. (Note: Unordered list content)

©2013-2025 Tencent Cloud. All rights reserved.

8. Click **Save** to complete the modification, and the database will be audited according to the new audit rules.

©2013–2025 Tencent Cloud. All rights reserved.

Modifying Audit Log Retention Period

Last updated: 2025-02-08 10:58:52

Overview

TencentDB for MongoDB audit service supports long-term storage of audit logs based on business needs to meet database compliance regulatory requirements. You can adjust the retention period of audit logs at any time according to the security compliance requirements of business audit rules.

Prerequisites

- Activated a TencentDB for MongoDB instance with database audit enabled.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Database Audit**.
3. Above the **Database Audit** page on the right, select the region.
4. In the top-right corner of the audit instance list, select an instance whose **Audit Status** is **Enabled**.
5. Click the instance ID with auditing enabled to go to the **Audit Log** page and view the corresponding logs.
6. In the top-right corner of the **Audit Log** page, click **Service Settings**.

7. In the pop-up dialog box, modify the log retention period and click **Submit**.

> **Note:**
> Note: To meet the security compliance requirements for the retention period of audit logs, we recommend users select a retention period of 180 days or above.

©2013-2025 Tencent Cloud. All rights reserved.

Shutting Down Database Audit Service

Last updated：2025-02-08 10:59:13

Overview

When your database instances no longer need auditing, please shut down the service promptly to avoid unnecessary charges.

Prerequisites

- Activated a TencentDB for MongoDB instance with database audit enabled.
- The TencentDB for MongoDB replica set or sharding instance is in **Running** status.

Operation Steps

1. Log in to the TencentDB for MongoDB console.
2. In the left navigation bar, select **MongoDB** > **Database Audit**.
3. Above the **Database Audit** page on the right, select the region.
4. In the top-right corner of the audit instance list, select an instance whose **Audit Status** is **Enabled**.
5. Click the instance ID with auditing enabled to go to the **Audit Log** page and view the corresponding logs.
6. In the top-right corner of the **Audit Log** page, click **Service Settings**.

7. In the pop-up dialog box, select **Shut Down the Service** after **Log Retention Period**, and click **Submit**.

> **Notes:**
> To meet the security compliance requirements for the retention period of audit logs, we recommend users select a retention period of 180 days or above.

©2013-2025 Tencent Cloud. All rights reserved.

Data Migration Guide
Creating a DTS Migration Task
Migration Overview

Last updated: 2025-02-08 11:00:01

Migration Services

Tencent Cloud Data Transmission Service (DTS) is a database data transfer service that integrates data migration, synchronization, and subscription, helping users easily complete database migration to cloud without service interruption. DTS for MongoDB can migrate data to the cloud database in one go, supporting the migration of both full and incremental data, i.e., migrating the historical data of the source database before migration and the newly written data during migration.

Migration Overview

Content Overview	Note:
Supporting Capabilities	Introduce the current features and use cases of using DTS for MongoDB.
Usage Instructions	Introduce the precautions and related restrictions during the use of DTS for MongoDB.
Migration Operation Guide	Introduce the specific operation steps for using DTS for migration.
Pre-validation Failure Handling	It is crucial to perform necessary verification before starting the DTS migration task to prevent and resolve database connection issues, table conflicts, and other potential issues. This article introduces common issues in the verification process.

©2013-2025 Tencent Cloud. All rights reserved.

Supported Capabilities

Last updated: 2025-02-08 11:00:23

Supported Scenarios and Versions

support migration between different architectures, replica set > replica set/sharded cluster, sharded cluster > replica set/sharded cluster, single node > replica set/sharded cluster

Source Database	Target Database	Scenario Description
Self-Built database (IDC self-built, CVM-based) on MongoDB 2.6, 2.8, 3.0, 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, 5.0, 6.0, 7.0	TencentDB for MongoDB 4.0, 4.2, 4.4, 5.0, 6.0, 7.0	Migration of on-premises databases to the cloud.
Databases in other clouds on MongoDB 2.6, 2.8, 3.0, 3.2, 3.4, 3.6, 4.0, 4.2, 4.4, 5.0, 6.0, 7.0 AWS MongoDB Atlas 6.0, 7.0	TencentDB for MongoDB 4.0, 4.2, 4.4, 5.0, 6.0, 7.0	Other cloud vendors migrating databases to Tencent Cloud database.
TencentDB for MongoDB 3.2, 3.6, 4.0, 4.2, 4.4, 5.0, 6.0, 7.0	TencentDB for MongoDB 4.0, 4.2, 4.4, 5.0, 6.0, 7.0	• Migration within the same Tencent Cloud region or between different regions. • Database migration between the same main accounts of Tencent Cloud, database migration between different main accounts. • Migration between different versions of Tencent Cloud MongoDB instances • Migration between Tencent Cloud MongoDB replica set and sharded cluster

Features

Feature Category	Feature Subitem or Description	Supported Capabilities
Migration Object	-	Database, Collection
Migration Type	-	• Full Migration • Full + Incremental Migration
Key Task Management Operations	Retry	Supported
	Create a Similar Task	Supported
Incremental synchronization	DML Synchronization (INSERT/UPDATE/DELETE)	Supported
	DDL Synchronization	INDEX: createIndexes, createIndex, dropIndex, dropIndexes COLLECTION: createCollection, drop, collMod, renameCollection, convertToCapped DATABASE: dropDatabase, copyDatabase Supports DDL operations for Replica Set and Sharding Cluster.
Consistency Check	Check Object	All Migration Objects/From Definition Object
	Verification Method	Row Count Comparison/Content Verification/Sampling Comparison

©2013-2025 Tencent Cloud. All rights reserved.

Usage Instructions

Last updated: 2025-02-08 11:00:42

Source Database Impact

During full data synchronization, DTS consumes certain source database resources, which may increase the load and pressure on the source database. If your database configuration is low, we recommend performing data migration during off-peak periods.

Target Database Impact

During the migration process, DTS uses the system service account to create a table under the `TencetDTSData` database on the target end with a task ID (e.g., table name `dts-xxxxx`) to record CHECKPOINT. This facilitates breakpoint resume in case of task interruption.

Migration Architecture

1. The related information for shard migration is as follows:
   1.1 Before shard cluster migration, it is recommended to clean up orphaned documents in the source cluster. Otherwise, this may lead to data validation inconsistency after migration. For information on how to clean up orphaned documents, please refer to the MongoDB official documentation cleanupOrphaned.
   1.2 During shard migration, do not enable sharding on the migrating database tables at the source end to avoid data distribution inconsistencies between the source and target ends. If sharding is enabled on the migrating database tables at the source end during migration, check the shard status on the target end. If sharding is not enabled on the target end, manually enable it. For specific operations on enabling sharding, please refer to the MongoDB official documentation Shard a Collection.
   1.3 The source end is a sharded cluster of TencentDB for MongoDB version 3.2. During migration, all shard keys are treated as hash shard keys by default. If you want to use range shard keys on the target end, create range shard keys on the target end before data migration.
2. Incremental migration is not supported for self-built single-node instances as they have no oplog.

Notes

1. Do not perform the following operations during the migration process, as they will cause the migration task to fail.
   - Do not modify or delete user information (including username, password, and permissions) in the source and target databases and port numbers.
   - Do not perform oplog cleanup operations on the source database.
   - During the data migration phase, do not delete the target database TencetDTSData.
2. During the data migration phase, operate on the target data with caution to avoid inconsistent final data.

©2013-2025 Tencent Cloud. All rights reserved.

Migration Operation Guide

Last updated: 2025-02-08 11:01:08

Overview

Based on DTS, MongoDB data migration supports full and incremental data migration. This means both historical data in the source database before the migration and new data written during the migration can be migrated.
This document provides operation guidance on migrating data from MongoDB to TencentDB for MongoDB using the DTS data migration feature.

Preparations

1. Carefully read the Usage Instructions to understand the feature constraints and precautions.
2. Please establish the access channel between DTS and the database according to your access type in advance. For details, refer to Network Preparations.
   - IDC self-built database/other cloud vendors' databases: Access methods can be "Public Network/Direct Connect/VPN Access/Cloud Connect Network".
   - Self-built databases on CVM: Access method selection is "Self-built on cloud host".
   - <TencentDB instance: Access method selection is "Cloud Database".
3. It is recommended to create a read-only account for migration in the source database. Refer to the method below.

```
# Syntax example when the source database is a replica set or single node
use admin
db.createUser({user: "username",pwd: "password",roles:[{role: "readAnyDatabase", db: "admin"},{role:
"read", db: "local"}]})

# Syntax example when the source database is a sharded cluster
use admin
db.createUser({user: "username",pwd: "password",roles:[{role: "readAnyDatabase", db: "admin"},{role:
"read", db: "local"},{role: "read", db: "config"}]})
```

4. The target database is a Tencent Cloud database instance. You can use mongouser for migration or create an account by yourself. The reference method for creating an account by yourself is as follows.

```
db.createUser({user:"username",pwd:"password",roles:[{role:"readWriteAnyDatabase",db:"admin"}]})
```

Operation Steps

1. Log in to the DTS Console, select **Data Migration** from the left navigation, click **Create Migration Task**, and enter the new migration task page.
2. On the Create Migration Task page, select the source instance type and region, target instance type and region, specifications, etc., then click **Purchase Now**.

Parameter Configuration	Description
Creation Model	• Create a new task: create a brand new task. • Create a similar task: quickly create a task with the same configuration as a historical task. The options such as database type, access method, billing mode, and migration type in the new task are pre-filled to match the historical task. Users can also modify them as needed.
Source Instance Type	Please select according to your source database type, which cannot be changed after purchase. In this scenario, select "MongoDB".
Source Instance Region	Select the source database region. If the source database is a self-built one, select a region nearest to it.

©2013-2025 Tencent Cloud. All rights reserved.

Target Instance Type	Please select according to your target database type, which cannot be changed after purchase. In this scenario, select "MongoDB".
Target Instance Region	Select the target database region.
Version	The default is NewDTS, no modification is needed.
Specification	Currently, only Medium specification is supported.
Task Name	• Select **name after creation**, the default task name is the same as the task ID. After the migration task is created, you can rename the task. • Select **Name Now** and enter the task name in the input box below.

3. After completing the purchase, the page will automatically jump to the data migration task list. Please select the task you just purchased for configuration.

   If you have purchased multiple regions or are configuring cross-region tasks, the task list will display according to the region of the target instance. You can switch regions at the top to find the purchased tasks.

4. On the source and target database settings page, complete the task settings, source database settings, and target database settings.

> **Note:**
> • Enter the previously created read-only account for the source database; otherwise, the pre-check step will fail.
> • If the source or target database is a Tencent Cloud database instance, DTS will use a system service account to export/write data during migration. For example, if the source database is a Tencent Cloud database instance, DTS will use the read-only account provided by the user to connect to the source database and also use a system service account to export data from the source database; if the target database is a Tencent Cloud database instance, DTS will use the account provided by the user to connect to the target database and also use a system service account to write data to the target database.

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Configuration	Description
Task Name	Set a business-significant name for easy task identification.
Execution Mode	• Immediate Execution: The task will start immediately after the pre-check passes. • Scheduled Execution: Set a start time for the task. The task will not start after the pre-check; it will start at the scheduled time.
Access Type	Please choose based on your scenario. For preparations of different access types, please refer to Overview. • Public Network: The source database can be accessed through the public network. • Self-built Cloud Host: The source database is deployed on a Tencent Cloud CVM. • Direct Connect: The source database can be connected to the Tencent Cloud VPC through Direct Connect. • VPN Access: The source database can connect to Tencent Cloud VPC through VPN Connections. • Cloud Database: The source database is a Tencent Cloud database instance. • CCN: The source database can connect to Tencent Cloud VPC through CCN. If the source database is an IDC self-built database or other cloud vendors' databases, the access type can be "Public Network/Direct Connect/VPN Access/CCN". If the source database is a self-built database on

©2013-2025 Tencent Cloud. All rights reserved.

	CVM, select "CVM-based Self-created". If the source database is a Tencent Cloud database instance, select "Cloud Database".
Architecture	**Access Type** When "CVM-based Self-created/Direct Connect/VPN Access/CCN" is selected, this parameter is displayed. Please choose according to the actual situation. • Replica Set: Refers to the source database being a replica set architecture, consisting of a Primary node and one or more Secondary nodes. • Cluster Migration: Refers to the source database being a sharded cluster architecture, consisting of mongos nodes, config servers, shard nodes, and other components. • Single Node: Refers to the source database cluster being a single node for read and write operations. **Note:** After choosing an architecture type and passing the connectivity test, it cannot be changed to another architecture type, or the task will error.
Cluster Migration	**Architecture** When "Cluster Migration" is selected, the following parameters need to be configured. • Node – mongod: Please enter the mongod node IP and port, or domain name and port. For multiple nodes, enter one per line; only one mongod per shard is needed. Example: 1xx.xx.55.77:6xx9 • Node – mongos: Please enter the mongos node IP and port, or domain name and port. • Node – Config Server: Please enter the Config Server node's IP and port, or domain name and port.
Public network	When **Access Type** is "Public network", configure the following parameters. • Host Address: IP address or domain name of the source database. • Port: The port used by the source database.
Self-built on cloud hosts	**Access Type** When "Self-built on cloud hosts" is selected, the following parameters need to be configured. • Cloud Host Instance: Instance ID of the CVM. • Port: The port used by the source database.
Direct Connect	**Access Type** When "Direct Connect" is selected, the following parameters need to be configured. • VPC Direct Connect Gateway: Only VPC Direct Connect Gateway is supported for Direct Connect. Please confirm the network type associated with the gateway. • VPC: Select VPC and subnet. • Host Address: IP address or domain name of the source database. • Port: The port used by the source database.
VPN Access	**Access Type** When "VPN Access" is selected, the following parameters need to be configured. • VPN Gateway: VPN Gateway, please select the VPN gateway instance accessed through the VPN Gateway. • VPC: Select VPC and subnet. • Host Address: IP address or domain name of the source database. • Port: The port used by the source database.
CloudDB	**Access Type** When "CDB" is selected, the following parameters need to be configured. Cloud Database Instance: Select the source database instance ID.

Cloud Connect Network	**Access Type** When "CCN" is selected, the following parameters need to be configured.
	When accessing via CCN, both same-account CCN and cross-account CCN are supported. Due to the extensive network configuration required, please refer to Migrating Self-built Databases to Tencent Cloud Databases via CCN.
	• Host Network Environment: Please select according to the actual situation. For example, if the source database is a Tencent Cloud database instance, select "Tencent Cloud"; if the source database is a self-built IDC database, select "Self-built IDC"; if the source database is from other cloud vendors, select the corresponding network.
	• Host Address: IP address or domain name of the source database.
	• Port: The port used by the source database.
	• Account of the CCN instance
	  ○ My account: The CCN resources and DTS belong to the same Tencent Cloud root account.
	  ○ Other account: The CCN resources and DTS belong to different Tencent Cloud root accounts.

©2013-2025 Tencent Cloud. All rights reserved.

	• Private network CCN: CCN instance name. • Accessed VPC: Select the accessed VPC and subnet. The accessed VPC refers to the VPC in CCN over which the DTS migration linkage is connected. Please select a VPC other than the source database's VPC from all VPCs associated with the CCN. • Accessed VPC region: The region of the source database selected during task purchase must be the same as the region of the accessed VPC; otherwise, DTS will change the former to the latter.

Authentication Required	Do you need to authenticate the security of the username and password in the source database? Select "Required" if you need to fill in the following parameters.
	• Authentication repository: The name of the database that requires authentication, which is the database name owned by the account performing the migration task. Only "admin" is supported.
	• Authentication Mechanism: Currently only SCRAM-SHA-1 is supported.
	• Account and Password Selection
	  ○ Same Account and Password: Select this parameter if the same account and password are used for mongod, mongos, and Config Server roles, and enter the unified account and password.
	  ○ Different Account and Password: Select this parameter if different accounts and passwords are used for mongod, mongos, and Config Server roles, and enter the account and password for each separately for mongod, mongos, and Config Server.

Account/Password	Account/Password: The account and password of the source database.
Connection Method	• Non-encrypted connection: When the source end is non-AMS Mongo Atlas, only non-encrypted connection can be selected. • Mongo Atlas SSL: When the source end is AWS Mongo Atlas, you can choose whether to use SSL encrypted connection as needed. For scenarios with high data security requirements, it is recommended to select Mongo Atlas SSL.

5. Test the connectivity of the source and target instances. If the connectivity test failed, please refer to Failed Connectivity Test for handling.

6. On the Set migration options and select migration objects page, set migration options and migration objects.

©2013-2025 Tencent Cloud. All rights reserved.

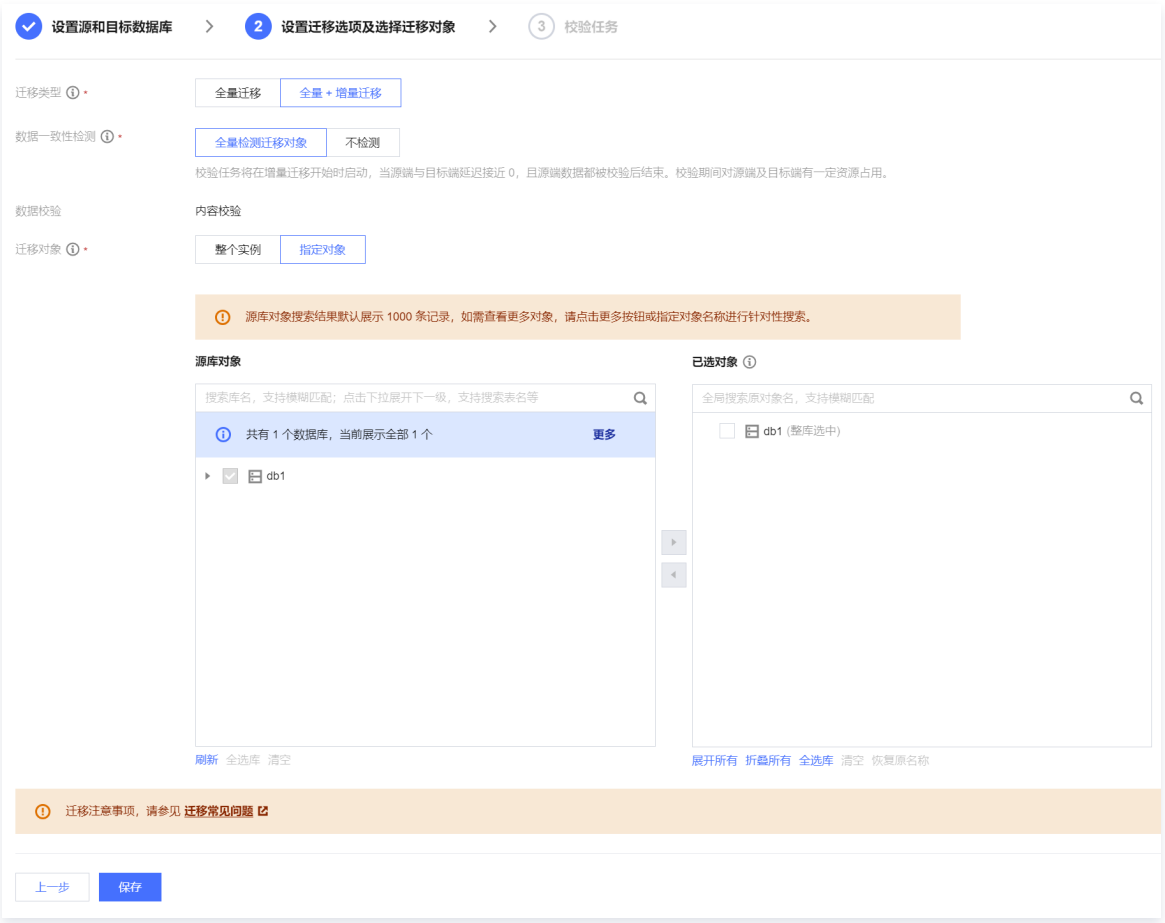

Parameter Configuration	Description
Migration Type	Please choose according to your scenario. • Full migration: The entire database will be migrated. The migrated data will only be existing content of the source database when the task is initiated but not include the incremental data written to the source database after the task is initiated. • Full + Incremental migration: The migrated data will 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.
Data Consistency Check	The consistency check is supported only when the **Migration Type** is configured as "Full + Incremental Migration" to compare data in detail between the source and target databases after migration. When the **Migration Type** is configured as "Full Migration", the check is not performed. • Full Check for Migration Objects: Once the incremental sync is completed, DTS will automatically trigger a full consistency check task. • Do not check: No data validation will be performed. If needed, users can manually trigger it after incremental synchronization is completed. For details, refer to Creating Data Consistency Check Task.
Data verification	When the **data consistency check** option is set to "Full scan migration object", the default type of consistency check displayed is "Content check".
Migration Object	Entire instance: Migrate the entire instance, but exclude system databases, such as system objects in Postgres, while migrating roles and User Metadata Definition. Specified objects: Migrate specified objects.
Specified objects	Select the objects to be migrated in Source Database Object and move them to the Selected Object box.

©2013–2025 Tencent Cloud. All rights reserved.

7. On the validation task page, complete the pre-migration validation work, and click **Start Task**. If the validation task failed, you can refer to Pre-validation Failure Handling to fix the problem and re-initiate the validation task.
   - Failed: It indicates that a check item failed and the task is blocked. You need to fix the problem and run the validation task again.
   - Alarm: It indicates that a check item doesn't completely meet the requirements, and the task can be continued, but the business will be affected. You need to assess whether to ignore the alarm or fix the problem and continue the task based on the alarm message.

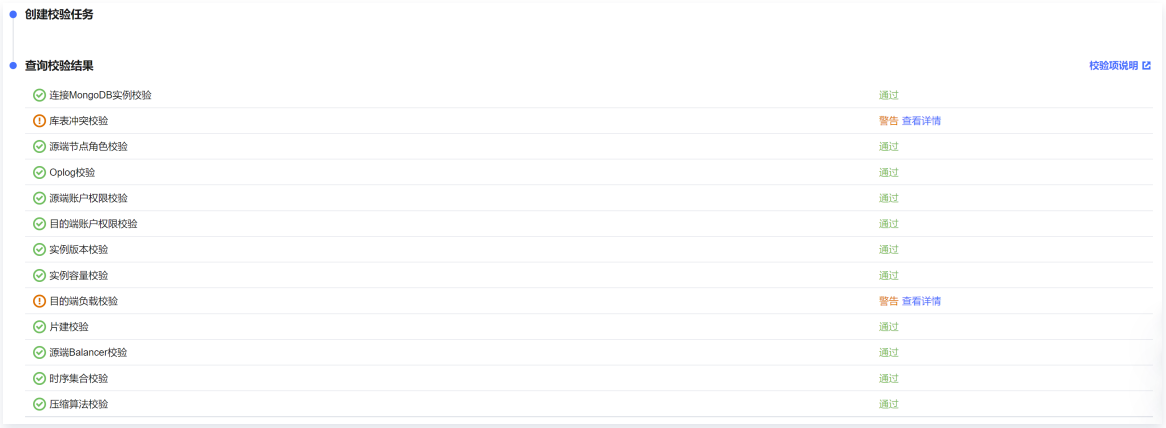

8. Return to the Migration Task List, and the synchronization task is Running.

   If you need to view task details, delete tasks, or perform other operations, please click the corresponding task operation. For more details, refer to Task Management. If a task error occurs, please refer to Error Handling.

9. Stop the task.
   - Selecting **Full migration**: The task will automatically end upon completion, manual termination is not required.
   - Selecting **Full + Incremental migration**: After the full migration is completed, it will enter the incremental data synchronization phase, which does not end automatically and requires the user to manually end the task.

     Please, when the incremental synchronization is completed (i.e., the status is "Preparation complete"), and there is zero-second delay between the target and the source databases, click **Complete** in the **Action** column to end the migration task.

10. (Optional) If a cutover is needed, after stopping the task and the task status becomes **Task successful**, you can proceed with the formal cutover. For more details, refer to Cutover Description.

©2013-2025 Tencent Cloud. All rights reserved.

Pre-validation failed processing Verify Connecting to TencentDB for MongoDB Instance

Last updated: 2025-02-08 11:01:46

Check details

The source and target databases need to be properly connected. If not, a connection failed error will occur.

Causes of Issues

- The source database's network or server has a security group or firewall configured.
- The source database has restricted the source IP address.
- The network port is not opened.
- The database account or password is incorrect.

Fix

Please address the issue according to the corresponding cause mentioned.

©2013-2025 Tencent Cloud. All rights reserved.

Table conflict verification

Last updated: 2025-02-08 11:02:11

Check Requirements

In MongoDB migration scenarios, the target instance can have the same name as the source database, but it must not contain any data (only empty tables are allowed).

Fix

If there is a conflict error, delete the corresponding database and table in the target database, or delete the data inside the duplicate name database in the target database.

©2013-2025 Tencent Cloud. All rights reserved.

Source Database Node Role Check

Last updated: 2025-02-08 11:02:35

Check details

- Check requirements: For MongoDB migration tasks, when the source end is sharded, information for the corresponding mongos, config server, and mongod nodes must be provided.
- Check explanation: The information for mongos, config server, and mongod nodes must not be mixed up; otherwise, it will lead to data migration disorder, such as entering mongos node information in the mongod input field. Note, each shard only needs one mongod node.

Fix

- Enter the correct node information in the DTS task configuration items.
- Only one mongod node is required for each shard.

©2013-2025 Tencent Cloud. All rights reserved.

Oplog Check

Last updated: 2025-02-08 11:02:56

Check details

- Check requirements: oplogs can be obtained from the source database during full + incremental migration.
- Check instructions: Incremental migration needs to replay the Oplog. If the source end's local library does not have oplog.rs or an oplog.$main table, the Oplog cannot be obtained.

Fix

Start the source end as a Replica Set or in Master-slave mode to ensure that operations generate Oplog and it is recorded in the source end's local library.

©2013-2025 Tencent Cloud. All rights reserved.

Verify source end or target end account permissions

Last updated: 2025-02-08 11:03:21

Check details

Check whether the user has operational permissions for the database. Specifically refer to the corresponding document below.
Permission requirements for data migration: MongoDB Data Migration

Fix

If the user does not have operational permissions, get authorized based on the permission requirements in the check details and re-execute the verification task.

©2013-2025 Tencent Cloud. All rights reserved.

Instance version verification

Last updated: 2025-02-08 11:03:42

The source and target database versions must be supported by MongoDB.

©2013-2025 Tencent Cloud. All rights reserved.

Instance Capacity Verification

Last updated: 2025-02-08 11:04:03

Check Requirements

In the MongoDB migration scenario, the storage space of the target database must be at least 1.3 times the size of the tables to be migrated in the source database.

Fix

- Delete some data from the target database to free up sufficient space.
- Upgrade the storage specifications of the target database and use an instance with larger capacity for migration.

©2013-2025 Tencent Cloud. All rights reserved.

Target Database Load Check

Last updated: 2025-02-08 11:04:24

Check details

- Check requirements: DTS migration will increase the load in the target database. If there is a business running in the target database during migration, a verification warning will be triggered. It will not block the task but will affect the business. You need to assess and determine whether to ignore the warning.
- Impact on business: MongoDB DTS uses logical sync for data migration, which will cause certain pressure on the CPU load of the target database. If there is a business running in the target database, you need to assess and initiate the migration task with caution.

Modification Method

Stop any business running in the target database and run the verification task again.

©2013-2025 Tencent Cloud. All rights reserved.

Shard Key Check

Last updated: 2025-02-08 11:04:44

Check details

- Check requirement: if the target database is a sharded instance, you can preset the shard key in it. If the table shard keys in the source and target databases are different, a warning will be triggered. It will not block the task but will affect the business. You need to assess and determine whether to ignore the warning.
- Impact on the business: in some cases where shard keys are inconsistent, the migration or sync task will fail.

Fix

If you preset the shard key in the target, use the following command to perform the sharding operation on the source.

```
sh.shardCollection("<database>.<collection>", { <shard key> : "hashed" } , false, {numInitialChunks:
number of preset chunks})
```

Run the verification task again.

©2013-2025 Tencent Cloud. All rights reserved.

Source Database Balancer Check

Last updated: 2025-02-08 11:05:06

Check details

- Check requirements: if the source database is a sharded instance, you need to disable the balancer in it before you can start migration.
- Inspection Note: Incremental Migration will fetch Oplog. With the Balancer enabled, moveChunk on the source might lead to data inconsistency at the destination.

Fix

1. Log in to the source database.
2. Use the following command to disable the source-side Balancer.

```
sh.stopBalancer()
sh.getBalancerState()
```

3. Run the verification task again.

©2013-2025 Tencent Cloud. All rights reserved.

Pipeline Legality Validation

Last updated: 2025-02-08 11:05:27

In MongoDB data subscription, DTS will validate the source database according to the native pipeline legality requirements of Change Stream. For syntactic errors unrelated to user data, an interface prompt will be provided. Please modify according to the prompt.

©2013-2025 Tencent Cloud. All rights reserved.

Time Series Collection Validation

Last updated: 2025-02-08 11:06:02

Verification Details

MongoDB 5.0 and above support Time Series Collections. When migrating from version 5.0 and above to a lower version, this validation item will fail if the source database contains Time Series Collections.

Fix

In the scenario of migrating from version 5.0 and above to a lower version, only select Non-time Series Collections when configuring tasks and choosing migration objects.

©2013-2025 Tencent Cloud. All rights reserved.

Compression Algorithm Verification

Last updated: 2025-02-08 11:06:26

Verification Details

Verify if the compression algorithm used by the source database matches with the target database. If they are different, a warning will be generated. The warning will not block the migration, and users can ignore it to continue the task.
Note that when verifying the compression algorithm used by the target database, an arbitrary system table is used. After modifying the compression algorithm, the system table's compression algorithm will not change, so there may be inaccurate warnings. If it is confirmed that the target database is already using the new compression algorithm, the warning can be ignored.

Fix

The same data occupies different disk sizes under different compression algorithms. If users want the target database to use the same compression algorithm as the source database, please modify the compression algorithm of the target database.

©2013-2025 Tencent Cloud. All rights reserved.

Consistency Check After Migration
Consistency Check Feature Description

Last updated: 2025-02-08 11:07:17

Feature Description

Data Consistency Check, namely DTS compares the table data between the source database and the target database during data migration and provides the comparison result and inconsistency details to help users quickly determine the business cutover time.

Notes

1. Data consistency check compares only the objects selected in the source database and objects migrated to the target database. If you write data into the target database during migration, then the written data will not be included in the consistency check.
2. A data consistency check task may increase the load in the source database instance. Therefore, you need to perform such tasks during off-peak hours.
3. Tasks for data consistency verification can be created and executed multiple times, but only one verification task can be "running" at any given moment. That is, the next verification task can only start after the previous one has ended or been terminated.
4. If a user chooses to **complete** or **terminate** a DTS task before the data consistency verification task has ended, the data consistency verification task will fail.
5. When creating a consistency verification, the system will automatically create the `dts_verify_result` database on the target end to record content related to consistency verification. The tables created under the `dts_verify_result` database are styled as follows:
   - diff_5xxxxxxxx4231: Stores the detected inconsistent data
   - diff_meta_5xxxxxxxxx4231: Stores the detected inconsistent metadata
   - result_5xxxxxxxxx4231: Records the result after each stage of checking is completed
   - status_5xxxxxxxxx4231: Records the verification progress
6. When the DTS link's target end is in the São Paulo Region, consistency check is temporarily not supported.

Restrictions

Currently, check tasks are imperceptible to the DDL operations. If you perform DDL operations in the source database during migration, the check result will be inconsistent with the actual data, and you need to initiate another check task to get the accurate comparison result.

Check Scheme

DTS verifies and compares data, including data from full migration and additional incremental migration data from the source end during the migration. The full data check compares the data from the source and destination row by row. Once the incremental data check thread finds that the full data comparison is completed, it immediately enters the incremental data check stage and gets the timestamp when the full data started (incremental logs are synced at the beginning of full data). It loops to get the incremental operation logs (Oplogs) from the source end and compares the differences between the source and destination. The comparison ends and the verification results are output when the delay between the data of the destination and the source end is less than 10 seconds.

©2013-2025 Tencent Cloud. All rights reserved.

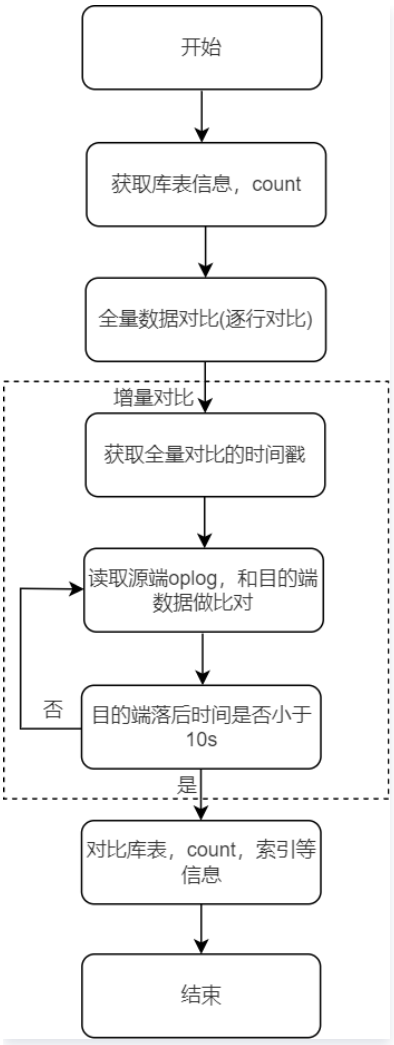

©2013–2025 Tencent Cloud. All rights reserved.

Creating Data Consistency Check Task

Last updated: 2025-02-08 11:07:37

Overview

Data consistency check tasks can be triggered automatically or created manually.

- Automatic trigger, when configuring the migration task, if **data consistency check** is selected for **full check of migration objects**, then a consistency validation task is automatically triggered when the task proceeds to the **sync increment** step.
- Manual creation, during the **sync increment** step of DTS tasks, manually create a consistency validation task, supporting multiple creations.

Triggering a data consistency check task automatically

In the data migration task **Set migration options and select migration objects** page, **data consistency check** is selected for **full check of migration objects**. This triggers a consistency validation task automatically when the task proceeds to the **incremental sync** step.

> **Notes:**
> Triggering a data consistency check task automatically: By default, all migration objects will be checked. If you need to screen check objects, manually create a data consistency check task.

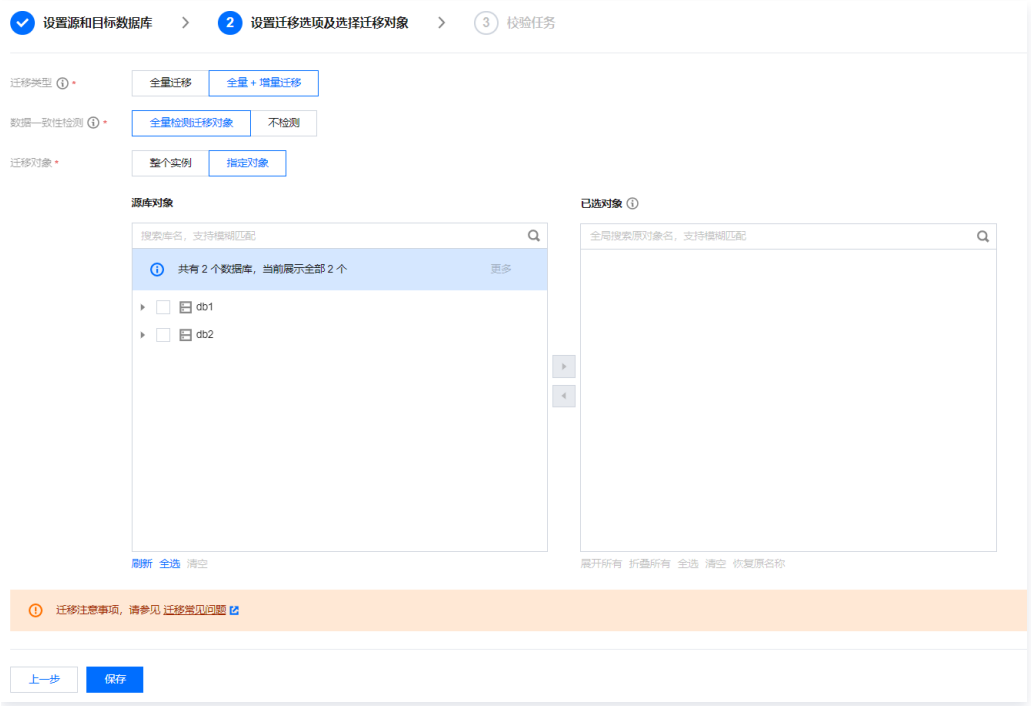

Creating a data consistency check task manually

1. Log in to the DTS console.
2. On the **data migration** page, select the migration task to be validated, and in the **action** column select **more** > **create data consistency check**.

3. Click **create data consistency check**.

> **Notes:**

©2013-2025 Tencent Cloud. All rights reserved.

> Data consistency validation needs to be created when the DTS task steps reach **sync increment**. If the interface button is grayed out, then the DTS task status does not meet the conditions, such as not reaching the **sync increment** step, task failure, or task termination.

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

5. After configuring the data consistency validation parameters, click **create and initiate consistency check**.

Parameters	Description
Task Name	Set the name of the check task.
Comparison type	Built-in Verification: The verification service is built into the DTS task, and only compares the data of migration objects.
Migration object mode	No configuration needed. This shows whether the migration task object selected in previous configurations is a whole instance or specified objects.
Check Object	• All migration objects: The verification range covers all selected migration objects in the migration task. • Custom Selection: Select the objects for verification from the selected migration objects.
Database Information	Supports verification of **indexes**, **shard keys**, and **database table information**. When both the source and target databases are sharded clusters, verification of **shard keys** is supported.
Data verification	• Row Count Comparison: Compare the row counts of the selected check objects. • Content Verification: Verify the content of the selected check objects. After selection, you can configure the sampling ratio.
Sampling Comparison	After selecting Content Verification, configure the sampling ratio, supporting 1%–100%.

©2013–2025 Tencent Cloud. All rights reserved.

	For some scenarios with large data volumes, full data verification may increase the load on the source database. Users should choose the sampling ratio based on their business needs.

View Consistency Verification Result

1. On the migration task homepage, in the **last verification result** column, you can view the results: consistent or inconsistent. Click **View More** to enter the **Data Consistency Verification** page.
2. In the data validation task list, select a specific task and click **Action** column's **View** to view the results of a single validation task.

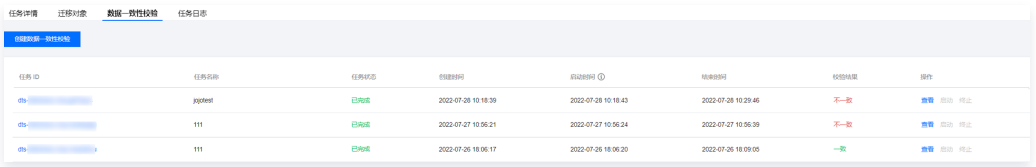

Example of consistent validation results:

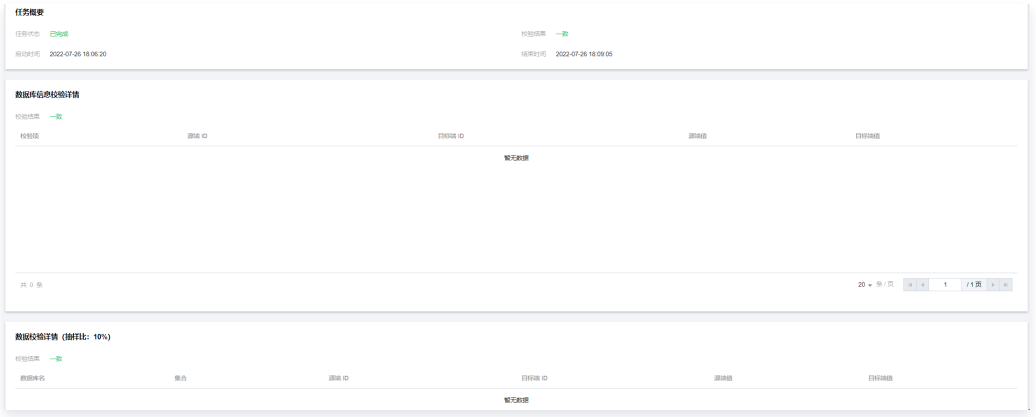

Example of inconsistent validation results:

> **Notes:**
> For inconsistent results, users need to manually confirm the corresponding data content of the source and target databases as prompted. For more handling methods, please refer to Common Consistency Verification Issues.

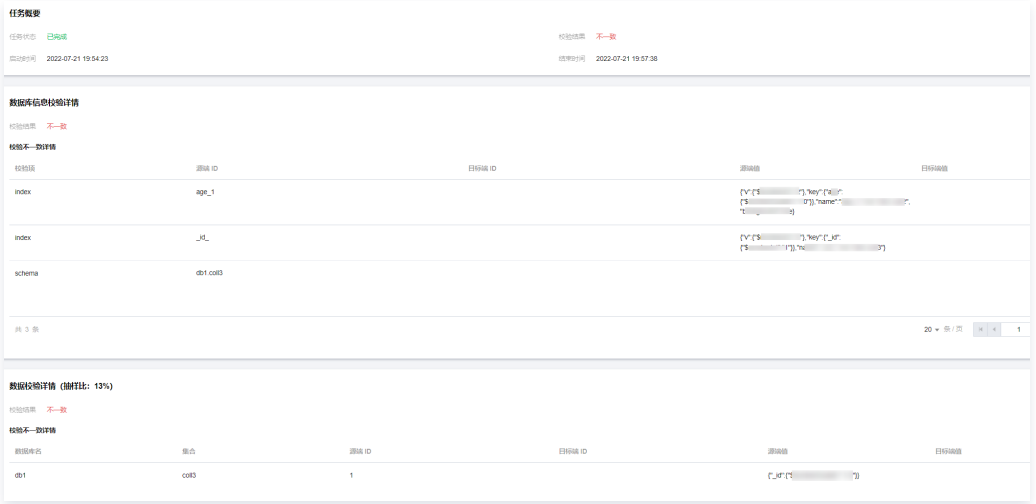

©2013–2025 Tencent Cloud. All rights reserved.

Common Consistency Check Issues

Last updated: 2025-02-08 11:07:58

Full data validation task is slow and time-consuming

Cause Analysis

The DTS consistency check policy compares full data in the source and target TencentDB for MongoDB databases row by row. If the data volume is large, the check task may take longer and use more system resources, thus affecting the source service and migration task.

Solution:

If the data volume is large and the check task is very slow, we recommend you stop the current task, manually create another consistency check task, and select **Row count check** as the data check method for quick comparison. Or, select **Content check** as the data check method and lower the proportion selected for **Sampling**, i.e., sampling a lower proportion of data for check, to improve verification efficiency. For detailed directions, see Creating Data Consistency Check.

- **Data Verification**
  - **Row count check**: Compares the number of rows of data in the source and target databases.
  - **Content check**: Compares the data content in the source and target databases row by row.
- **Sampling Comparison**: When selecting content check, it supports comparing a certain proportion of the data.

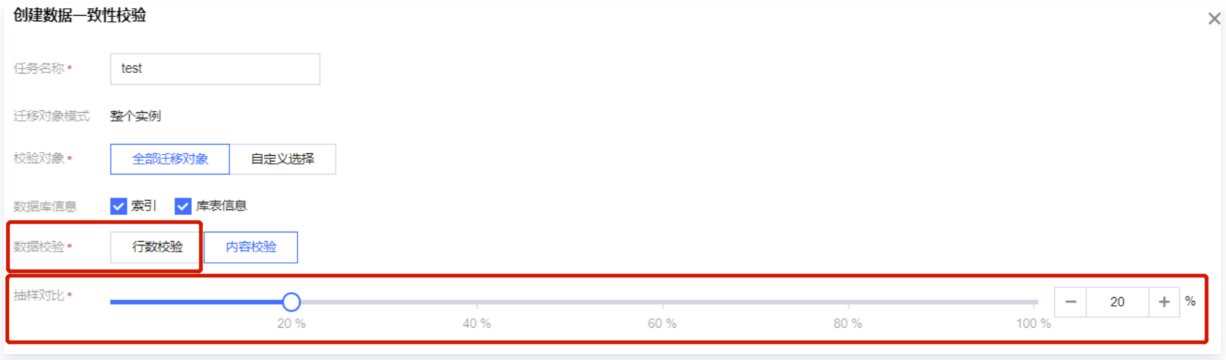

Data inconsistency

Data content inconsistency

Problem Phenomenon

Cause Analysis

During full data verification, data is continuously written to the source, and the Oplog is continuously generated. The incremental data verification task continuously reads the Oplog from the source. If the newly generated Oplog in the source has not reached the target's timestamp, there may be short-term data content inconsistencies, which is normal.

Solution

Judge the inconsistent data row by row. You can also initiate a new verification task to perform another manual verification. When the target catches up with the increment, the inconsistency will disappear.

Data row count inconsistency

Problem Phenomenon

©2013-2025 Tencent Cloud. All rights reserved.


数据行数校验详情

校验结果 不一致

校验不一致详情

数据库	集合	源端行数	目标端行数
testdb	coll	2	1

Reason One

During a full data check, data is continuously written to the source database, and the Oplog keeps being updated. The incremental data check task continuously reads the Oplog from the source database. If the new Oplog generated in the source database has not reached the timestamp in the target database, the number of data rows may be inconsistent for a short period of time, which is a normal phenomenon.

Reason Two

Solution

In such cases, you can use db.collection.countDocuments() for an accurate row count comparison. Note that this method scans the collection and poses some performance risks. For more information, please refer to db.collection.countDocuments() .

Index check

Problem Description

When creating a consistency check task, if the **database information** includes **index** comparison between the source and target, and you notice differences in the "v" and "background" fields between the source and target, but the check results do not indicate these inconsistencies.

Explanation of Reasons

TencentDB for MongoDB index verification policy ignores version information: differences in the "v" field and background creation: differences in the "background" field's content will not be displayed in the verification results.

©2013–2025 Tencent Cloud. All rights reserved.

Technical Scheme and Common Problems of Data Consistency Check

Last updated：2025-02-08 11:08:22

During data consistency check, DTS compares the collection data between the source and target databases and outputs the comparison result and inconsistency details for you to perform a business cutover stably and reliably.

Check Scheme

DTS checks and compares all the data migrated during full migration and incremental migration from the source database. A full data check compares the data in the source and target databases row by row. Once the thread of the incremental data check finds that the full data comparison is completed, it immediately starts the incremental data check to get the start timestamp of the full data check (to compare incremental data, it needs to start from the snapshot data at the start time of the full data), get the incremental oplog in the source database in a loop, and compare the differences between the source and target databases. When the time lag of data in the source and target databases is below 10 seconds, the comparison ends, and the check result is output.

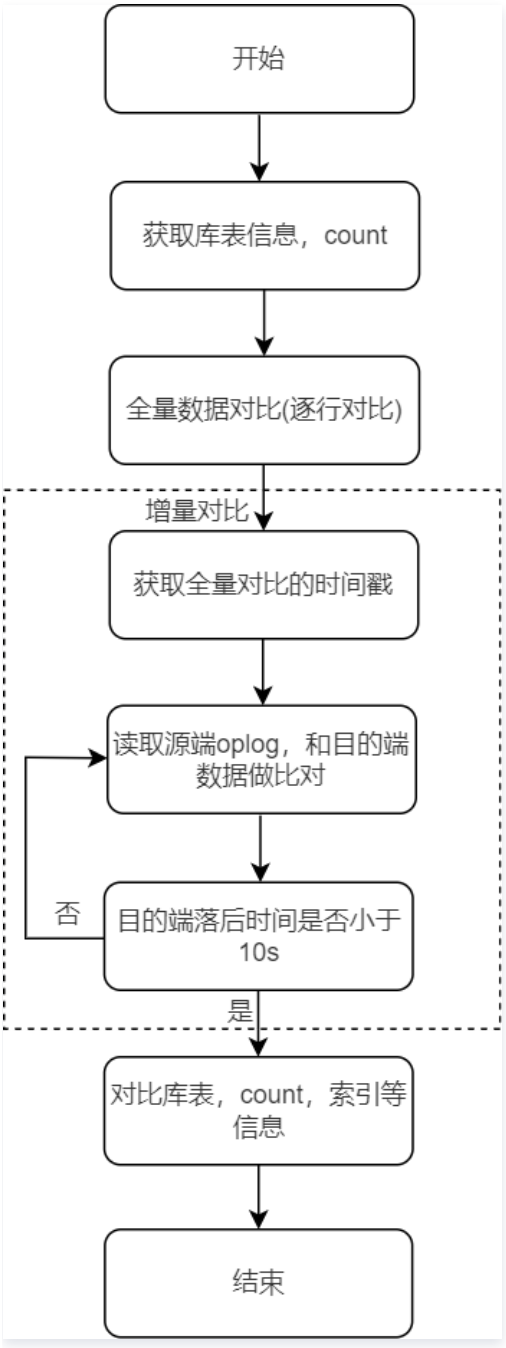

©2013-2025 Tencent Cloud. All rights reserved.

FAQs

Time-consuming full data check task

Cause Analysis

The DTS consistency check policy compares full data in the source and target TencentDB for MongoDB databases row by row. If the data volume is large, the check task may take longer and use more system resources, thus affecting the source service and migration task.

Solution:

If the data volume is large and the check task cannot be completed for a long time, we recommend stopping the current task, manually creating another consistency check task, and selecting **Row count check** as the data check method for quick comparison. Or, select **Content check** as the data check method and lower the proportion selected for **Sampling**, i.e., sampling a lower proportion of data for check, to improve verification efficiency. For detailed directions, see Creating Data Consistency Check Task.

- **Data Verification**
  - **Row count check**: Compares the number of rows of data in the source and target databases.
  - **Content check**: Compares the data content in the source and target databases row by row.
- **Sampling Comparison**: When selecting content check, it supports comparing a certain proportion of the data.

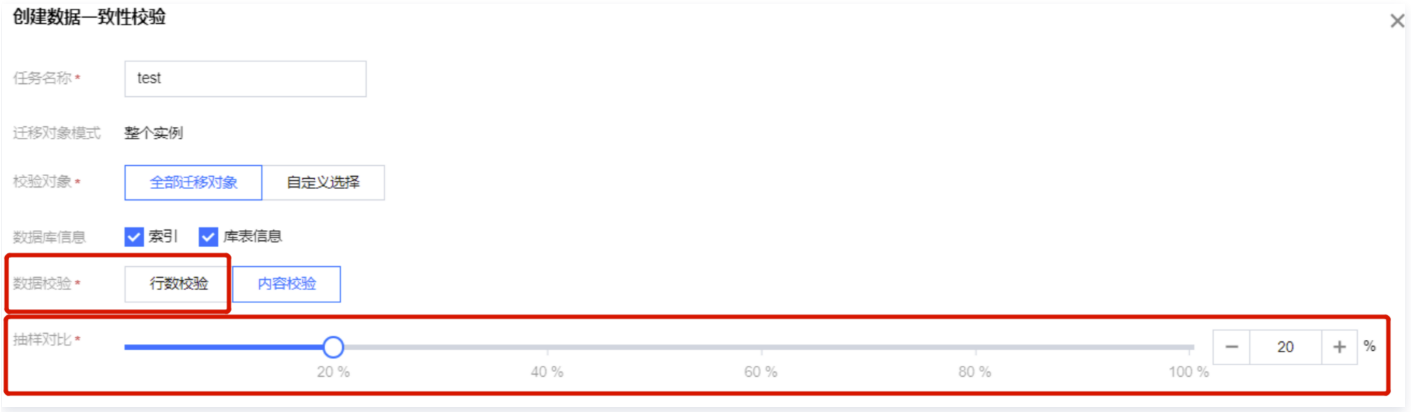

Data inconsistency

Data content inconsistency

Problem phenomenon

Cause Analysis

During full data verification, data is continuously written to the source, and the Oplog is continuously generated. The incremental data verification task continuously reads the Oplog from the source. If the newly generated Oplog in the source has not reached the target's timestamp, there may be short-term data content inconsistencies, which is normal.

Solution

Judge the inconsistent data row by row. You can also initiate a new verification task to perform another manual verification. When the target catches up with the increment, the inconsistency will disappear.

Data row count inconsistency

©2013-2025 Tencent Cloud. All rights reserved.

Problem phenomenon

Reason One

During a full data check, data is continuously written to the source database, and the Oplog keeps being updated. The incremental data check task continuously reads the Oplog from the source database. If the new Oplog generated in the source database has not reached the timestamp in the target database, the number of data rows may be inconsistent for a short period of time, which is a normal phenomenon.

Reason Two

Solution

In such cases, you can use db.collection.countDocuments() for an accurate row count comparison. Note that this method scans the collection and poses some performance risks. For more information, please refer to db.collection.countDocuments() .

Index check

Problem Description

When creating a consistency check task, if the **database information** includes **index** comparison between the source and target, and you notice differences in the "v" and "background" fields between the source and target, but the check results do not indicate these inconsistencies.

Explanation of Reasons

TencentDB for MongoDB index verification policy ignores version information: differences in the "v" field and background creation: differences in the "background" field's content will not be displayed in the verification results.

©2013–2025 Tencent Cloud. All rights reserved.

Creating MongoDB Data Subscription

Last updated: 2025-02-08 11:08:43

This document describes how to create a data subscription task in DTS for TencentDB for MongoDB.

Version Description

- Data subscription for TencentDB for MongoDB supports versions 3.6, 4.0, 4.2, and 4.4.
- TencentDB for MongoDB 3.6 only supports collection-level subscription.

Prerequisites

- You have prepared a TencentDB instance to be subscribed to, and the database version meets the requirements. For more information, see Databases Supported by Data Subscription.
- We recommend you create a read-only account in the source instance by referring to the following syntax. For console operations, see Create Read-only Account.

```
# Create an instance-level read-only account
use admin
db.createUser({
user: "username",
pwd: "password",
roles:[
       {role: "readAnyDatabase", db: "admin"}
      ]
})

# Create a read-only account for a specific database
use admin
db.createUser({
user: "username",
pwd: "password",
roles:[
       {role: "read",db: "Name of the specified database"}
      ]
})
```

Restrictions

- Currently, the subscribed message content is retained for 1 day by default. Once expired, the data will be cleared. Therefore, you need to consume the data promptly.
- The region where the data is consumed should be the same as that of the subscribed instance.
- The Kafka built in DTS has a certain upper limit for processing individual messages. When a single row of data in the source database exceeds 10 MB, this row may be discarded in the consumer.
- If the specified database or collection for the data subscription task is deleted in the source database, the subscription data (change stream) of the database or collection will be invalidated. Even if the database or collection is rebuilt in the source database, the subscription data cannot be resubscribed. You need to reset the subscription task and select the subscription object again.

SQL operations for subscription

Operation Type	Supported SQL Operations
DML	INSERT, UPDATE, DELETE
DDL	INDEX: createIndexes, createIndex, dropIndex, dropIndexes COLLECTION: createCollection, drop, collMod, renameCollection DATABASE: dropDatabase, copyDatabase

©2013-2025 Tencent Cloud. All rights reserved.

Subscription configuration steps

1. Log in to the DTS console, select **Data Subscription** in the left sidebar, and click **Create Subscription**.
2. On the **Data Transfer Service** page, refer to the table below to configure the interface parameters, and click **Purchase Now**.

Parameter Name	<0>Parameter meaning<0>	Configuration Method
Service Type	Select the type of Data Transmission Service. This document introduces the **Data Subscription** service.	Select **Data Subscription**.
Billing Mode	Select the billing method for this service. For billing details, see Data Transmission Service Billing Overview.	Both **Monthly Subscription** and **Pay-As-You-Go** are supported. For selection guidance, see Data Transmission Service Billing Mode.
Region	Select the region of the subscription service.	Keep it consistent with the region of the database instance to be subscribed to.
Database	Select the database type for the data subscription service.	Select **MongoDB**.
Version	Support consuming directly through Kafka client.	Select **Kafka Edition**.
Tag	Assign a tag to the data subscription service.	Click **Add**, and select **Tag Key** and **Tag Value** from the dropdown list.
Subscription Instance Name	Specify the method for setting the data subscription service name.	• **Name After Creation**: Set the name after creating the data subscription service. The default is name-*Subscription ID*. The Subscription ID is randomly assigned by the system. • **Immediate Naming**: Directly set the data subscription service name in the input box below.
Automatic Renewal	When selecting the **Billing Mode** as **annual and monthly subscription**, you need to set whether to enable auto-renewal. If the account balance is sufficient, the subscription service will be automatically renewed monthly after expiration.	It is recommended to check. If not, please pay attention to the service expiration warning information. For details, see overdue explanation.
Purchase Duration	When selecting the **Billing Mode** as **annual and monthly subscription**, you need to select the purchase duration of the service.	The longer the purchase duration, the more discounts you enjoy.
Purchase Quantity	Select the quantity of the service to purchase.	You can purchase up to 10 tasks in a single transaction.

3. After a successful purchase, return to the data subscription list and select the subscription task you just purchased. In the **Operation** column, click **Configure Subscription**.

4. On the **Select an Instance** tab of the **Subscription Configuration** wizard, configure the database information for the data subscription task and perform a connectivity test.

©2013-2025 Tencent Cloud. All rights reserved.

Parameter Name	<0>Parameter meaning<0>	Configuration Method
Subscription ID / Name	The ID and name of the subscription task. The task name defaults to name–Subscription ID.	Confirm the ID and name of the data subscription.
Instance Type	The default is **MongoDB**.	–
Region	The region where the subscription service is located.	Confirm the region.
Access Type	Select the type of source database to access the data subscription service.	Currently, only **CloudDB** is supported. It refers to TencentDB instances.
TencentDB instance	Select the specific MongoDB instance for the data subscription service.	In the dropdown list, select the specific instance id.
Account	Set the access account information for the MongoDB instance.	Enter the prepared read–only account information in the input box.
Password	Set the password for the access account of the MongoDB instance.	Enter the password information for the read–only account in the input box. Password–free access is not supported.
Number of Kafka Partitions	Select the number of Kafka partitions for the data subscription task. In Kafka, consumers can subscribe to one or more topics to obtain data, and then consume data from one or more partitions of each topic.	Supports selecting **1**, **4**, **8**. • A single partition can guarantee the order of messages, while multiple partitions cannot. If you have strict requirements on the order of consumed messages, set the number of Kafka partitions to 1. • Increasing the number of partitions can improve the throughput and parallelism of a Kafka cluster, as multiple consumers can consume different partitions simultaneously. However, increasing the number of partitions also increases the management and maintenance costs of the Kafka cluster and may cause data imbalance or delay issues.

©2013–2025 Tencent Cloud. All rights reserved.

Test Connectivity	Test the connectivity between the data subscription service and the TencentDB for MongoDB instance.	• Click **Test Connectivity** and wait for the connectivity test results. • If the test failed, troubleshoot according to the prompt method. You can click **Testing again** to recheck connectivity. The next step can only be performed if the test passes.

5. Click **Next,** on the **Subscription Type and Object Selection** page, configure the parameters as shown in the table below, and click **Save Configuration.**

Parameter Name	<0>Parameter meaning<0>	Configuration Method
Subscription ID / Name	The ID and name of the subscription task.	Confirm that the subscription task information is correct.
Subscription instance	The ID of the subscribed MongoDB database instance.	Confirm that the instance information is correct.
Data Subscription Type	Refers to the data type that subscribers can choose to subscribe to. For MongoDB databases, use the change stream feature to monitor changing data and achieve data subscription.	It is **Change Stream** by default and cannot be modified.
Subscription Object Level	The data subscription levels include full instance, database, and collection. • **Full instance**: Subscribe to full instance data. • **Database**: Subscribe to database-level data. After selection, **Task Configuration** allows only one database to be selected. • **Collection**: Subscribe to collection-level data. After selection, **Task Configuration** allows only one collection to be selected.	You can choose the data subscription level as needed. System databases such as admin, local, and config are not supported.
Task Configuration	When the **Subscription Object Level** is **Database** or **Collection**, this parameter is displayed to specify the database and collection.	• In the **select database tables** area, select one database or collection to subscribe to. • In the **Selected** area, confirm whether the selected database and collection are correct.
Output Aggregation Settings	This parameter configures whether to enable aggregation settings for the subscribed data before sending it to the subscriber.	• Enable: click to enable aggregation settings. Click **Add**, select an operator from the **Aggregation Operator** dropdown list, then enter an expression in the **Aggregation Expression** input box. • Click **Add** to add multiple aggregation expressions. The execution order of the aggregation pipeline follows the order of the added aggregation operations. For more information, see MongoDB Official Documentation.
Kafka Partitioning Policy	When the number of Kafka partitions in the previous step is not 1, you need to set a partitioning policy. • **By Collection Name**: Partitions the subscribed data from the source database by collection name. Data with the same collection name will be written to the same Kafka partition.	• Select **Custom Partitioning Policy**, click **Add** under **Custom Partitioning Policy**, and set the matching method for the database name or table name to be customized in the input boxes of **Database Name Matching Mode** and

©2013-2025 Tencent Cloud. All rights reserved.

	• **Custom Partitioning Policy**: Database and collection names of the subscribed data are matched through a regex first. Then, matched data is partitioned by collection name or collection name + objectid.	**Table Name Matching Mode** using regex. Then, in the **Partitioning Policy** dropdown list, select **By Collection Name** or **By Collection Name + objectid**. • Enabling the custom partitioning policy will prioritize matching the custom partitioning policy, followed by the Kafka partitioning policy. • For databases and tables that do not meet the above custom partitioning policy, they will be routed to Kafka partitions according to the default policy: "By Collection Name".
Custom Partitioning Policy	**Kafka Partitioning Policy** select **Custom Partitioning Policy** to display this parameter. Set the custom partitioning policy rules.	
Policy Combination Results	**Kafka Partitioning Policy** select **Custom Partitioning Policy** to display this parameter. Explain the combination results of the custom partitioning policy.	

6. On the **Pre-verification** page, the pre-verification task is expected to run for 2-3 minutes. After the pre-verification is passed, click **Start** to complete the configuration of the data subscription task.

> **Note:**
> If the verification fails, fix the problem as instructed in Fix for Verification Failure and initiate the verification again.

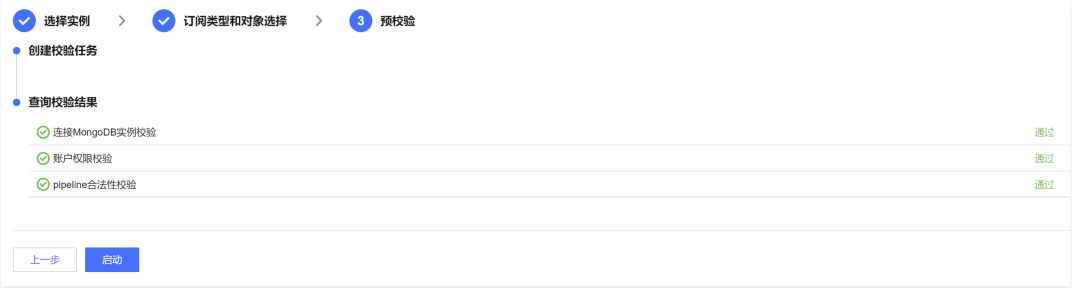

7. The subscription task will be initialized, which is expected to run for 3-4 minutes. After successful initialization, it will enter the **Running** status and start consuming data.

Subsequent Operations

1. Add a consumer group.
   The consumption of Data Subscription (Kafka Edition) depends on the consumer groups of Kafka; therefore, you must create a consumer group before data can be consumed. Data Subscription (Kafka Edition) allows you to create multiple consumer groups for multi-point consumption.
2. Data Consumption.
   After the subscription task enters the **Running** status, data consumption can begin. Kafka consumption requires password authentication. For multi-language sample code, refer to the demo in Data Consumption.

©2013-2025 Tencent Cloud. All rights reserved.