

Tencent Cloud TCHouse-D

Operation Guide

Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

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

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by the Tencent corporate group, including its parent, subsidiaries and affiliated companies, as the case may be. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

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

Contents

Operation Guide

Cluster Operation

- Creating and Terminating

- Scale-Out and Scale-In

- Configuration Adjustment

Monitoring and Alarm Configuration

- Cluster Monitoring

- FE Monitoring Metric

- BE Monitoring Metric

- Business Monitoring Metrics

- Alarm Configuration

Account Privilege Management

- Account Privilege Management

- Privilege Management Through a Client

Data Management

- Data Backup

- Data Recovery

- Cross-Cluster DataMigration

- Database Auditing

- Cold-Hot Data Tiering

Query Management

- Query Analysis

- Slow Query Analysis

Modify Configurations

- Modify Configurations Through Console

- Modify Configurations Through Client

Node Management

Log Analysis

SQL Studio

Enabling Resource Isolation

Operation Guide

Cluster Operation

Creating and Terminating

Last updated : 2025-01-06 09:44:57

Create Cluster

1. On the Tencent Cloud TCHouse-D page, click **Buy Now**, or log in to the [Tencent Cloud TCHouse-D Console](#) and click **Create Cluster** to proceed to the purchase page for configuration and purchase.
2. After completing the basic configuration, cluster configuration, and log configuration, the system will display the corresponding configuration costs. Click **Buy now** to create a cluster.

Cluster

Cluster name
6-36 characters; supports Chinese characters, letters, digits, hyphens (-), and underscores (_).

Kernel version 2.1 ▼

FE node type

Standard

Compute spec 4-core 16 GB ▼

Storage spec CLOUD_HSSD ▼ - 200 + GB
200 to 32000 GB for a single node

FE node count - 3 +

The number of nodes cannot exceed the result of the number of available subnet IP addresses in the selected network minus the number of BE nodes. If the number of subnet IP addresses is not enough, switch the subnet or VPC and try again.

BE node type

Standard

High-performance

Compute spec 4-core 16 GB ▼

Storage spec CLOUD_HSSD ▼ - 200 + GB
200 to 320000 GB for a single node

BE node count - 3 +

The number of nodes must be larger than or equal to 3 but cannot exceed the result of the number of available subnet IP addresses in the selected network minus the number of FE nodes. If the number of subnet IP addresses is not enough, switch the subnet or VPC and try again.

Configuration Item Description

Configuration Item	Description
Billing Mode	Yearly/monthly: Prepaid , i.e., you pay upfront when creating a cluster, then use the resources. Pay-as-you-go: Postpaid . You can request resources on-demand, and an expense bill is generated every hour. Charges are based on actual resource usage upon settlement.
Region	Currently, Tencent Cloud TCHouse-D supports regions such as Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, Nanjing, Hong Kong (China), Singapore, Frankfurt, Silicon

	Valley, and Tokyo. We recommend selecting a region closest to your business needs. Once a region is selected, it cannot be changed after purchase.
High Availability	<p>You can choose to enable the high-availability read-write pattern:</p> <p>In the non-high-availability pattern, only 1 FE will be deployed. It is not recommended for production environments, especially for online queries or real-time read-write scenarios</p> <p>In the high-availability read pattern, at least 3 FEs will be deployed. One FE node is a follower, and the other FE nodes are observers (high availability read).</p> <p>High-availability read-write: At least 5 FEs are deployed, of which 3 FE nodes are followers (high availability write), and the other FE nodes are observers (high availability read).</p>
AZ Deployment Mode (Cross-AZ high availability)	<p>Single AZ: All BE/FE nodes are deployed in one AZ</p> <p>Three AZs: It is optional when the number of availability zones in the current region ≥ 3. It is recommended to select the "read/write high availability" mode: at least 5 FE nodes, ensuring at least 1 FE follower node in each AZ</p> <p>BE nodes must be a multiple of 3, the replication_num of the table will be forcibly set to 3, and evenly distributed across each AZ to ensure data high availability</p>
Network	A VPC is an isolated, highly secure, dedicated network environment. You can create a VPC network and a subnet or select an existing one.
AZ and subnet	<p>Select AZs in different regions as needed on the purchase page.</p> <p>For three-AZ deployment, select 3 AZs and configure subnets</p>
Compute Node Type	<p>Compute nodes support three types:</p> <p>Standard: 4-core 16 GB, 8-core 32 GB, 16-core 64 GB, 32-core 128 GB, 64-core 256 GB, etc.</p> <p>Memory optimized: 4-core 32 GB, 8-core 64 GB, 16-core 128 GB, 32-core 128 GB, 64-core 256 GB</p> <p>High performance: 16-core 64 GB, 32-core 128 GB, 64-core 256 GB</p>
Table Name Case Sensitivity	<p>Table name case can't be modified after configuration, and it provides three patterns:</p> <p>Case Sensitivity</p> <p>Case-insensitive, stored as specified and compared in lowercase</p> <p>Case-insensitive, convert all table names to lowercase for storage</p>
Cloud Disk Encryption	<p>Cloud disk encryption is a free feature that effectively protects your data privacy, meets security compliance requirements, and requires no modification of your business code. Encryption can only be set at the time of purchase and is only supported for cloud disks, not local disks.</p> <p>There is no conversion between encrypted and non-encrypted cloud disks.</p> <p>After cloud disk encryption, an estimated performance loss of 10%-20% is expected.</p>

View Cluster Information

1. After the cluster is created, you can enter [Tencent Cloud TCHouse-D Console](#), select the region where the cluster is located, and view the status information of all the clusters in this region from the cluster list.

Resource ID/Name	Status (progress)	Health status	FE node	BE node	Kernel vers...	AZ	Network	Subnet	Billing mode	Creation time	Operatio
	Serving	Healthy	Standard, 3 node(s) 4-core 16GB, 200 GB	Standard, 3 node(s) 4-core 16 GB, 200 GB	2.1	ap-hongkong-2			Monthly subscription Expires on 2025-01-18 15:03:50	2024-11-18 15:03:50	Scale out Scale up Terminate More

2. In the cluster list, click **Cluster ID/Name** to view specific details of the cluster and perform cluster-level operations. In the cluster information page, you can view the basic information, cluster status, configuration information, and network information of the cluster.

You can modify the cluster name by clicking the edit icon next to **Cluster Name**.

Cluster Info

Cluster Monitoring

Accounts

Data Management

- Backup And Restore
- Cold-Hot Tiering
- Database Auditing

Queries

- Query Analysis
- SQL Studio

Configurations

- Modify Configurations
- Change History

Nodes

Log Analysis

Operation Logs

Lakehouse Integration

Cluster Info

Scale out Scale in Scale up/down

Basic info

Cluster ID: [redacted]

Cluster name: [redacted]

Billing mode: Monthly subscription

Creation time: 2024-11-18 15:03:50

Expiration time: 2025-01-18 15:03:50

Tag: [Change](#)

Configuration info

Kernel version: 2.1 (tencent-cdw-doris-2.1.7-3f71f0e-2410221)

High availability (HA): Read high availability

FE node spec: Standard, 4-core 16 GB / 3 node(s) / CLOUD_HSSD 200 GB

BE node spec: Standard, 4-core 16 GB / 3 node(s) / CLOUD_HSSD 200 GB

Default Alarm Policy: Enable

Table names are case-sensitive: Case-Sensitive

Cloud Block Storage Encryption: Close

Security Groups: None

Cluster status

Cluster status: Serving

Network info

AZ: ap-hongkong-2

VPC ID: [redacted]

Subnet ID: [redacted]

JDBC access address: 10.22.0.14:9030 [Copy connection string](#)

HTTP connection address: 10.22.0.14:8030 [Copy connection string](#)

Node info

No.	Node type	Node health	Node IP	Cluster VPC ID
1	FE(master)	Healthy	10.22.0.87	9.0.16.17
2	FE(observer)	Healthy	10.22.0.7	9.0.16.11
3	FE(observer)	Healthy	10.22.0.98	9.0.16.16
4	BE	Healthy	10.22.0.52	9.0.16.8
5	BE	Healthy	10.22.0.115	9.0.16.6
6	BE	Healthy	10.22.0.45	9.0.16.10

Total items: 6 10 / page 1 / 1 page

Terminating Cluster

Click **Operation > Terminate** in the cluster list to terminate a cluster as prompted. After termination, resources will be released and data will be cleared within 24 hours.

Terminate

1 Termination options > 2 Confirm

i • A terminated cluster will be retained as "Isolated" for 7 days. Please back up data in advance.

• If you terminate the cluster within 5 days after creation, the refund amount will be returned to your Tencent Cloud account unconditionally. If not within 5 days, the amount paid in cash or free credit will be returned to your Tencent Cloud account based on their percentage in the payment.

• If the cluster was purchased with a discount or voucher, the respective amount will **not be refunded**.

Cluster ID	Cluster name	Spec	Expiration time
		4-core 16 GB, 200 GB, 3 (FE) node(s)	2025-01-18
		4-core 16 GB, 200 GB, 3 (BE) node(s)	

[Next: Confirm](#) [Cancel](#)

Scale-Out and Scale-In

Last updated : 2024-12-03 11:30:32

Overview

Tencent Cloud TCHouse-D cluster provides a scale-in and scale-out feature, which helps you to scale the number of nodes in the cluster.

Scale-out refers to reducing the number of service nodes. If the current scale and performance of the cluster do not meet usage requirements, you can enhance the performance through scale-out.

Scale-in refers to reducing the number of service nodes. If the current usage of the cluster is low, and there is no significant demand for a long period, you can use scale-in to reduce cluster scale and save cost.

Before scale-in and scale-out operations, please ensure the following:

The cluster is in a stable operating state.

The main account is not in arrears and there are no outstanding orders.

Scale out

In the Scale-out process, the system can read and write, but there may be some fluctuations. The operation generally takes about 5-15 minutes. Please carry it out during non-peak business hours.

When the business is facing double growth in data storage volume and query volume, it is recommended to choose scale-out as a priority.

Note:

For a non-high-availability cluster (with only one FE node), expanding to 3 or more FE nodes will by default switch to a read high-availability pattern.

If read-write high availability is not specified, even when expanding to more than 5 FE nodes, it will remain in the read high-availability pattern.

Operation Guide

1. Sign in to [Tencent Cloud TCHouse-D Console](#), and choose **Operation > Scale-out** for the cluster you wish to operate.



Resource ID/Name	Status (progress)	Health status	FE node	BE node	Kernel version	AZ	Network	Subnet	Billing mode	Creation time	Operation
...	Serving	Healthy	Standard, 3 nodes 4-core 16GB, 200 GB	Standard, 3 nodes 4-core 16 GB, 200 GB	2.1	ap-hongkong-2	Pay-as-you-go	2024-11-25 15:22:46	Scale out Scale up/down Terminate More

2. In the cluster scale-out pop-up window, select the type and quantity of nodes to be scaled out, and you can calculate the change fee in real time. The standalone configuration of the new scale-out node is the same as the existing configuration of the cluster.

Cluster scale-out
✕

i

- Reads and writes on the cluster are still available during scale-out, but there may be some jitters. Scale-out will take about 5 to 15 minutes. We recommend you perform this operation during off-peak hours.
- Scale-out is recommended if both your data storage needs and queries increase.
- If a security group has been configured for the cluster, you cannot scale out the cluster in a new subnet. Make sure all ports of all IP addresses in the subnet are open before you scale out the cluster.

Cluster ID/name

Region/AZ HongKong/ap-hongkong-2

Billing mode Monthly subscription

High availability Read high availability

FE node spec Standard/4 cores 16 GB/3 nodes/CLOUD_HSSD200 GB

BE node spec Standard/4 cores 16 GB/3 nodes/CLOUD_HSSD200 GB

Node to scale FE node BE node

Change to Non-high availability Read high availability Read-Write high availability

-
3
+

You can increase the node count to 3–127. If available resources are insufficient, please [contact us](#).

For read high availability: The number of nodes must be greater than or equal to 3, with 1 node as a Follower, and the rest as Observers.
For read-write high availability: The number of nodes must be greater than or equal to 5, with 3 nodes as Followers and the remaining nodes as Observers.

Network

Subnet **61 subnet IPs in total,54available.**

If no IP address is available in the current subnet, you can switch the subnet.

Cost --

Confirm
Cancel

3. Click **Confirm**. After completing the payment, the cluster will start the scale-out process, which generally takes about 5-15 minutes.

4. After the scale-out is complete, go to the cluster details page where you can peek at the list of nodes after scale-out.

Scale-in

It is recommended to perform scale-in operations on one node at a time. During the scale-in process, system reading and writing can still be carried out, but there might be some fluctuations.

Based on the requirements for architectural stability, there are the following requirements for the number of target nodes after scale-in:

The number of FE nodes needs to be odd.

The number of FE and BE nodes in a high-availability pattern cluster should be equal to or greater than 3.

Minimum quantity requirements for node types: At least 1 FE node and at least 3 BE nodes are required.

To ensure the safety of data and a reasonable number of replicas, before scale-in, the storage data of the node to be removed needs to be migrated to the normal node. The specific time of scale-in execution depends on the size of the data, ranging from a few minutes to several hours. Please carry it out during non-peak business hours.

Operation Guide

1. Sign in to [Tencent Cloud TCHouse-D Console](#), choose **Operation > Scale-in** in the cluster you want to operate.

Resource ID/Name	Status (progress)	Health status	FE node	BE node	Kernel version	AZ	Network	Subnet	Billing mode	Creation time	Operation
[Icon]	Serving	Healthy	Standard, 3 node(s) 4-core 16GB, 200 GB	Standard, 3 node(s) 4-core 16 GB, 200 GB	2.1	ap-hongkong-2	[Progress]	[Progress]	Pay-as-you-go	2024-11-25 15:22:46	Scale out Scale up/down Terminate More
[Icon]	Serving	Healthy	Standard, 3 node(s) 4-core 16GB, 200 GB	Standard, 3 node(s) 4-core 16 GB, 200 GB	2.1	ap-hongkong-2	[Progress]	[Progress]	Monthly subscription Expires on 2024-12-22 17:27:21	2024-11-22 17:27:21	Scale out Scale up/down Scale in Create log service Terminate More
[Icon]	Serving	Status warning	Standard, 3 node(s) 4-core 16GB, 200 GB	Standard, 3 node(s) 4-core 16 GB, 200 GB	2.1	ap-hongkong-2	[Progress]	[Progress]	Monthly subscription Expires on 2024-12-22 17:17:43	2024-11-22 17:17:43	Cross-Specific Scale out Scale up/down Scale in Close Default Alia Terminate more

2. On the cluster scale-in page, select the node to be scaled in; the fee will be displayed in real-time.

3. Click **Confirm**. The cluster will begin the data migration and scale-in operations.

Cluster scale-in ✕

- Scale-in can be performed on only a node at a time.
- Removing multiple BE nodes at a time may cause data loss or take a longer time. We recommend you remove them one by one.
- Reads and writes on the cluster are still available during scale-in, but there may be some jitters.

Cluster ID/name: [ID] AZ: ap-hongkong-2

Billing mode: Pay-as-you-go High availability Read high availability

Kernel version: 2.1

FE node spec: Standard/4 cores 16 GB/3 nodes/CLOUD_HSSD200 GB

BE node spec: Standard/4 cores 16 GB/3 nodes/CLOUD_SSD200 GB

Node to scale: FE node BE node

Change to: Non-high availability Read high availability Read-Write high availability

For non-high availability: The number of nodes must be equal to 1 and as a Follower node.
For read high availability: The number of nodes must be greater than or equal to 3, with 1 node as a Follower, and the rest as Observers.
For read-write high availability: The number of nodes must be greater than or equal to 5, with 3 nodes as Followers and the remaining nodes as Observers.

Nodes to remove

Available: 2

Node IP	AZ	Node status
<input type="checkbox"/> 10.22.0.57(master)	ap-hongko...	Running
<input checked="" type="checkbox"/> 10.22.0.120(observer)	ap-hongko...	Running
<input checked="" type="checkbox"/> 10.22.0.66(observer)	ap-hongko...	Running

Selected: 2

Node IP	AZ	Node st...
10.22.0.120(observer)	ap-hong...	Running
10.22.0.66(observer)	ap-hong...	Running

Target node count: 1

New cost: [Bar chart]

Confirm
Cancel

4. After the scale-in is complete, go to the cluster details page, where you can view the list of nodes after scale-in.

Configuration Adjustment

Last updated : 2025-01-16 17:56:09

You can use the configuration adjustment feature to adjust the specifications of all FE/BE nodes within the cluster. The compute nodes will restart all nodes of the same type, making the system unreadable and unwritable. Please proceed during off-peak business hours. Upgrading the storage configurations will not affect the cluster's usage.

Note:

Standard nodes support both computing configuration adjustment and scaling up of storage configurations.

High-performance nodes do not support scale-up and scale-down in either compute or storage specifications.

Operation Guide

1. Log in [Tencent Cloud TCHouse-D Console](#), and select **Operation > Scale up/down** in the cluster where the operation is needed.



Resource ID/Name	Status (progress)	Health status	FE node	BE node	Kernel version	AZ	Network	Subnet	Billing mode	Creation time	Operation
...	Scaling	Healthy	Standard, 3 node(s) 4-core 16GB, 200 GB	Standard, 3 node(s) 4-core 16 GB, 200 GB	2.1	ap-hongkong-2	vpc-...	subnet-...	Pay-as-you-go	2024-11-25 15:22:46	Scale up Scale down Terminate More

2. In the Scale up/down pop-up for the cluster, select the type of node, computing configuration, and storage configuration to adjust.

When scaling up, the computing configuration and storage configuration of nodes can be independently or simultaneously scaled up.

For scaling down, the compute configuration of nodes can only be separately scaled down.

If the scaling involves a change in compute configuration, the relevant nodes need to be restarted. You can choose either rolling restart or full restart as needed. If only the storage configuration is changed, no restart is required.

Monitoring and Alarm Configuration

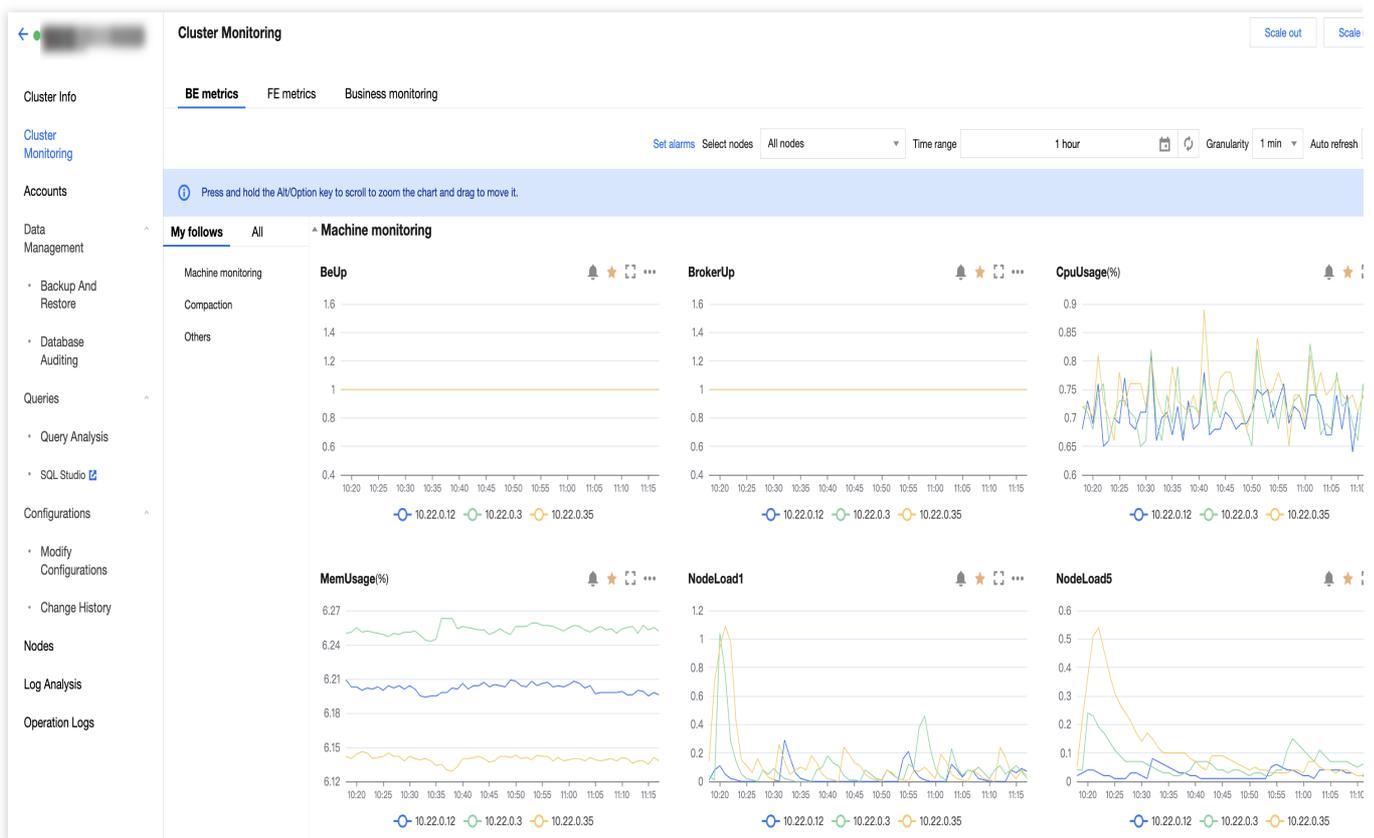
Cluster Monitoring

Last updated : 2024-06-27 10:46:35

Tencent Cloud TCHouse-D provides a wide range of monitoring metrics, allowing you to grasp the operating status of the cluster, and configure alarms based on metrics to pay attention to operation information in real time and achieve quick responses.

Monitoring Overview

Log in to [Tencent Cloud TCHouse-D Console](#), click on cluster **ID/Name** to enter the cluster detail page, and click on **Cluster Monitoring** to check various performance metrics of the cluster.



Note

You can view and export the data for specific monitoring metrics.

Monitoring collection interval is 10 seconds. The timing granularity exhibited can be selected by the user, and the smallest granularity is 1 minute. The monitoring data within the timing granularity are aggregated and shown. The

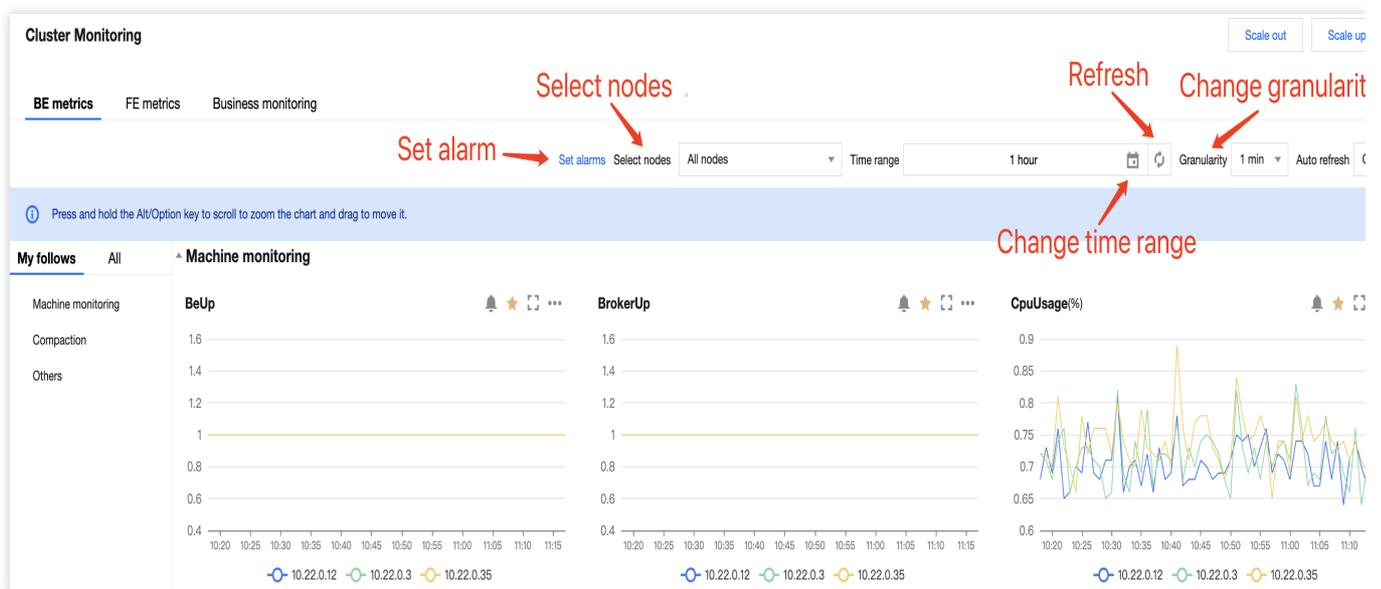
aggregation logic of metric data is to take the maximum.

Feature Details

In general, cluster monitoring is divided into three modules: BE metric monitoring, FE metric monitoring, and business monitoring.

You can adjust the time range of monitoring data to view historical data in different time periods. You can also adjust the timing granularity of monitoring data to view data on different dimension levels.

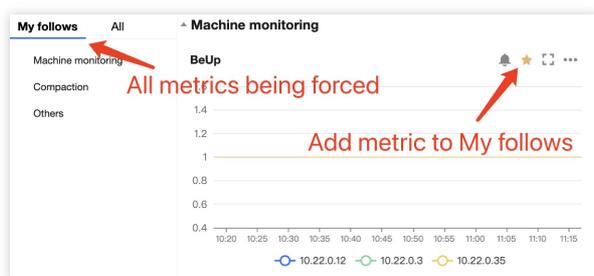
The monitoring page supports auto-refresh, which updates monitoring charts automatically according to the refresh interval selected by the user, achieving real-time data viewing.



You can add metrics that you wish to focus on to **My Follows** to quickly view them.

Add metrics to My follows: Click the star-button in the metric chart or the star-button next to the corresponding metric in the metric list.

View favorite metrics: You can switch to **My Follows** in the metric list and view all important metrics being focused.



Tencent Cloud TCHouse-D Service Monitoring

Cluster monitoring data will be reported to [Tencent Cloud Observability Platform](#), where you can retrieve the monitoring data.

FE Monitoring Metric

Last updated : 2025-03-07 14:49:18

On the FE metric page, you can view various FE metrics, as shown in the table below.

Category	Metric Name	Unit	Metric Alarm Name
Machine Monitoring	FeUp	-	fe_up
	BrokerUp	-	broker_up
	CpuUsage	%	cpu_usage
	MemUsage	%	mem_usage
	DiskUsage	%	disk_usage
	NodeLoad1	-	node_load1
	NodeLoad5	-	node_load5
	NodeLoad15	-	node_load15
	NodeNetworkReceiveBytesTotal	Bytes/s	node_network_receive_bytes_total
	NodeNetworkTransmitBytesTotal	Bytes/s	node_network_transmit_bytes_total
	NodeDiskReadsCompletedTotal	Count/s	node_disk_reads_completed_total
	NodeDiskWritesCompletedTotal	Count/s	node_disk_writes_completed_total
	SystemDiskUsage	%	system_disk_usage
Query	Added SQL Cache Num	Count	cache_added_sql
	Hit SQL Cache Num	Count	cache_hit_sql
	MySQL Port Connection Num	Count	connection_total
	Query Qps	Count/s	doris_fe_qps
	Query Error Num	Count	query_err
	Query Slow Num	Count	query_slow
	Querying Fragment Instance Num	Count	doris_fe_query_instance_num

	MySQL Port Rps	Count/s	doris_fe_rps
	Query Time Quantile 75	MilliSecond	doris_fe_query_latency_ms__0_75
	Query Time Quantile 95	MilliSecond	query_latency_ms_0_95
	Query Time Quantile 99	MilliSecond	query_latency_ms_0_99
	Query Time Quantile 999	MilliSecond	query_latency_ms_0_999
	Query Hive Table Num	Count	query_hive_table
Load	Finished Insert Num	Count	job_loadinsertfinished
	Cancelled Insert Num	Count	job_loadinsertcancelled
	Finished Broker Load Num	Count	job_loadbrokerfinished
	Cancelled Broker Load Num	Count	job_loadbrokercancelled
	Finished Delete Job Num	Count	job_loaddeletefinished
	Cancelled Delete Job Num	Count	job_loaddeletecancelled
	Scheduling Routine Load Num	Count	doris_fe_job__loadroutine_loadneed_schedu
	Running Routine Load Num	Count	doris_fe_job__loadroutine_loadrunning
	Paused Routine Load Num	Count	doris_fe_job__loadroutine_loadpaused
	Stopped Routine Load Num	Count	doris_fe_job__loadroutine_loadstopped
	Cancelled Routine Load Num	Count	doris_fe_job__loadroutine_loadcancelled
	Running Rollup Job Num	Count	doris_fe_job__alterrolluprunning
	Running Schema Change Num	Count	doris_fe_job__alterschema_changerunning
	Routine Load Error Rows	Count	routine_load_error_rows
	Routine Load Size	Byte	routine_load_receive_bytes
	Routine Load Rows	Count	routine_load_rows
	Tablet Num	Count	doris_fe_tablet_num

	Tablet Max Compaction Score	Count	doris_fe_tablet_max_compaction_score	
Transaction	Tablet Compaction Max Score	Count	doris_fe_max_tablet_compaction_score	
	Publishing Transaction Num	Count	doris_fe_publish_txn_num	
	Report Task Queue Size	Count	doris_fe_report_queue_size	
	Scheduled Tablet Num	Count	doris_fe_scheduled_tablet_num	
	Unhealthy Tablet Num	Count	tablet_status_count_unhealthy	
	Rejected Transaction Num	Count	txn_counter_reject	
	Began Transaction Num	Count	txn_counter_begin	
	Success Transaction Num	Count	txn_counter_success	
	Failed Transaction Num	Count	txn_counter_failed	
	Running Transaction Num	Count	doris_fe_txn_num	
	Cancelled Insert Job Num	Count	job_loadinsert_jobcancelled	
	Finished Insert Job Num	Count	job_loadinsert_jobfinished	
	JVM	JVM Max Heap Size	Byte	jvm_heap_size_bytes__max
		JVM Used Heap Size	Byte	jvm_heap_size_bytes__used
JVM Used Non Heap Size		Byte	jvm_non_heap_size_bytes__used	
GC Young Num		Count	jvm_young_gc_count	
GC Young Time		MilliSecond	jvm_young_gc_time	
GC Old Num		Count	jvm_old_gc_count	

	GC Old Time	MilliSecond	jvm_old_gc_time
	JVM Thread Num	Count	jvm_thread__count
	JVM Thread Peak Num	Count	jvm_thread__peak_count
	JVM Blocked Thread Num	Count	jvm_thread__blocked_count
	JVM Waiting Thread Num	Count	jvm_thread__waiting_count
Meta	Editlog Write Num	Count	edit_log_write
	Editlog Read Num	Count	edit_log_read
	Editlog Clean Success Num	Count	edit_log_clean_success
	Editlog Clean Failed Num	Count	edit_log_clean_failed
	Editlog Write Time Quantile 95	MilliSecond	doris_fe_editlog_write_latency_ms__0_95
	Editlog Write Time Quantile 999	MilliSecond	doris_fe_editlog_write_latency_ms__0_999

BE Monitoring Metric

Last updated : 2025-03-07 14:49:18

On the BE metric page, you can view various BE metrics, as shown in the table below.

Category	Metric Name	Unit	Metric Alarm Name
Machine Monitoring	BeUp	-	be_up
	BrokerUp	-	broker_up
	CpuUsage	%	cpu_usage
	MemUsage	%	mem_usage
	DiskUsage	%	disk_usage
	NodeLoad1	-	node_load1
	NodeLoad5	-	node_load5
	NodeLoad15	-	node_load15
	NodeNetworkReceiveBytesTotal	Bytes/s	node_network_receive_bytes_total
	NodeNetworkTransmitBytesTotal	Bytes/s	node_network_transmit_bytes_total
	NodeDiskReadsCompletedTotal	Count/s	node_disk_reads_completed_total
	NodeDiskWritesCompletedTotal	Count/s	node_disk_writes_completed_total
	SystemDiskUsage	%	system_disk_usage
Compaction	Cumulative Compaction Max Score	Count	doris_be_tablet_cumulative_max_compaction_score
	Base Compaction Max Score	Count	doris_be_tablet_base_max_compaction_score
	Waiting Compaction Permits	Count	doris_be_compaction_waiting_permits
	Used Compaction Permits	Count	compaction_used_permits
	Cumulative Compaction Merged Size	Byte	compaction_bytes_total_cumulative
	Base Compaction Merged Size	Byte	compaction_bytes_total_base
Load	Tablet Writer Num	Count	tablet_writer_count

	Stream Load Request Num	Count	streaming_load_requests_total
	Stream Load Duration	MilliSecond	streaming_load_duration_ms
	Stream Load Current Processing Num	Count	doris_be_streaming_load_current_process
	Stream Load Rollback Transaction Num	Count	stream_load_txn_request_rollback
	Routing Load Task Num	Count	doris_be_routine_load_task_count
	Tablet Sink Load Size	Byte	load_bytes
	Stream Load Received Size	Byte	stream_load_receive_bytes
	Stream Load Loaded Rows	Count	stream_load_load_rows
	Tablet Sink Sent Rows	Count	load_rows
Query	Canceled Fragment Num	Count	timeout_canceled_fragment_count
	Segment Read Rows	Count	segment_read_segment_row_total
	Segment Read Num	Count	segment_read_segment_read_total
	Query Scan Rows	Count	query_scan_rows
	Query Scan Size	Byte	query_scan_bytes
	Fragment Thread Pool Queue Size	Count	doris_be_fragment_thread_pool_queue_si
	Fragment Request Num	Count	fragment_requests_total
	Fragment Request Duration	MicroSecond	fragment_request_duration_us
	Fragment Instance Num	Count	doris_be_plan_fragment_count
	Fragment Instance Num	Count	fragment_instance_count

Others	Segment Meta Cache Capacity	Byte	doris_be_cache_capacity__segmentmetac
	Hit Ratio of File Cache	%	file_cache_hits_ratio
	Chunk Allocator Reserved Size	Byte	doris_be_chunk_pool_reserved_bytes
	Unused Rowsets Num	Count	unused_rowsets_count
	File Descriptor Num in Use	Count	process_fd_num_used
	Total Rowset Num	Count	doris_be_all_rowsets_num
	Uploaded Rowset Size	Byte	upload_total_byte
	Memory Allocated Size	Byte	doris_be_memory_allocated_bytes
	Total Segment Num	Count	doris_be_all_segments_num
	Remote Scan Thread Queue Size	Count	doris_be_remote_scan_thread_pool_queu
	Thread Num in Process	Count	doris_be_process_thread_num
	Local Scan Thread Num	Count	doris_be_local_scan_thread_pool_thread_
	Local Scan Thread Queue Size	Count	doris_be_local_scan_thread_pool_queue_

Trash Data Size	Byte	doris_be_disks_trash_used_capacity__d
Remote Data Size	Byte	disks_remote_used_capacity
Local Data Size	Byte	disks_local_used_capacity
JVM Thread Num	Count	jvm_thread_count

Business Monitoring Metrics

Last updated : 2024-12-03 11:30:32

On the business monitoring page, you can view various business metrics, such as slow query statistics, as shown in the table below.

Metric Name	Unit	Metric Name
DorisSlowQueryRecordsCount	Count	doris_slow_query_records_count
DorisUnhealthyTableCount	Count	doris_unhealthy_tablet_count

You can identify situations where the monitoring statistics query time exceeds the threshold, carry out corresponding troubleshooting and optimization, and improve system performance.

Note:

When configuring monitoring alarms for business monitoring metrics, you need to select "TCHouse-D/FE ALARM" in the policy type. The reason is that such metric monitoring actually depends on the FE node.

1 Configure Alarm Policy > 2 Configure Alarm Notification

Basic Info

Policy Name

Remarks

Configure Alarm Rule

Monitoring Type

Policy Type **TCHouse-D / FE ALARM** 3 exist. You can create 297 more static threshold policiesThe current account has 0 policies for dynamic alarm thresholds, and 20 more policies can be created.

Tag

[+ Add](#)

Alarm Object

Trigger Condition

Metric Alarm

When meeting of the following metric conditions, the metric will trigger an alarm. Enable alarm level feature.

Threshold Type

If (statistical p... > 0 % at 1 consecuti... then Alarm once an h...

[Add Metric](#)

Previous step

Next step: Configure Alarm Notification

Alarm Configuration

Last updated : 2024-12-03 11:30:32

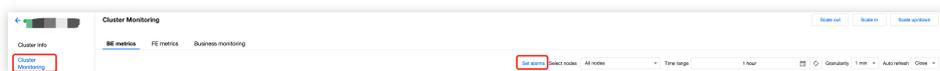
Tencent Cloud supports the creation of comprehensive and flexible alarm policies. All monitoring metrics can be set to raise alarms for any anomalies. The monitoring object can be all the clusters under your management, or a specific cluster or node.

Alarm rules support single or multiple combined metrics. Each metric can define alarm conditions for different time granularities.

Tencent Cloud supports self-defined alarm notification recipients and methods, and also allows the saving of a type of alarm notification as a template for reuse in multiple alarm policies.

Operation Guide

1. Enter the cluster monitoring page, and click **Set alarms** to create a new alarm policy.



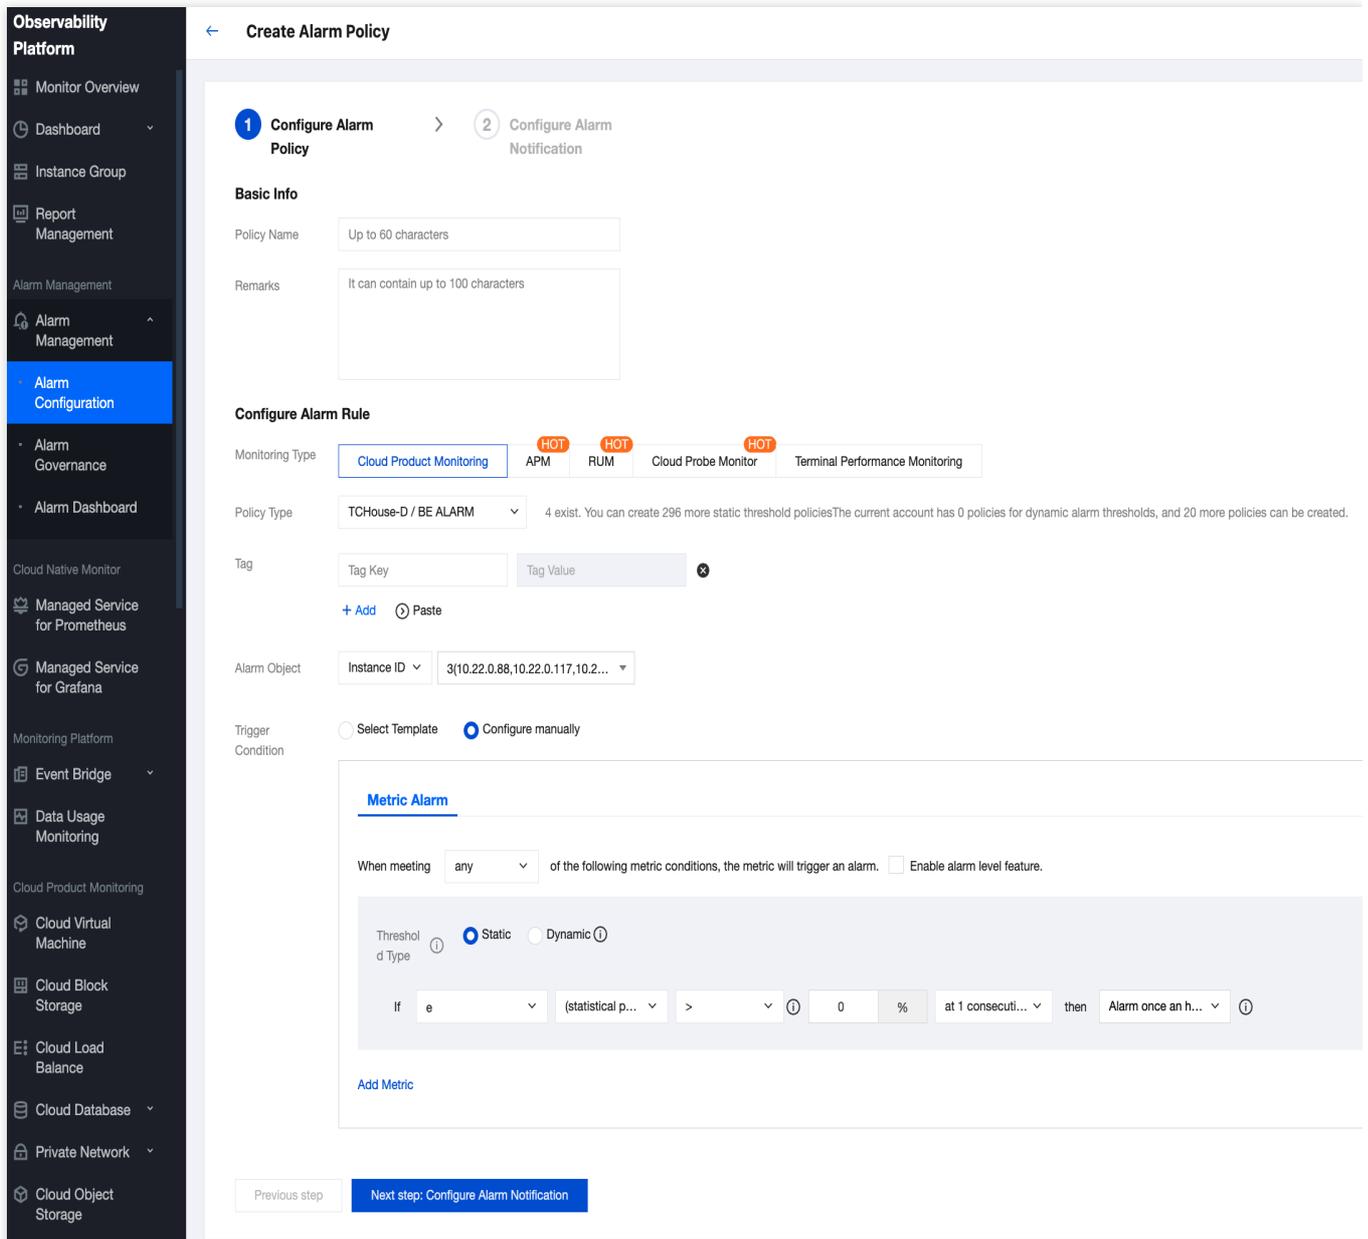
2. In the newly opened **Create Alarm Policy** window, fill in the basic alarm information as per the table below, filter the instance object (cluster name) requiring configuration, select metric items, and configure the monitoring threshold.

Setting Option	Note
Policy Name	Name of the alarm policy being configured
Monitoring Type	Default Tencent Cloud Product Monitoring
Policy Type	Default CDWDORIS/BE Alarms or CDWDORIS/FE Alarms
Alarm Object	Select a cluster or node using the instance ID
Trigger condition	Configure alarm policies and related thresholds for metrics
Notification Template	Choose an existing template or create a new notification template

3. After configuration, click **Finish** to submit the alarm policy. For more information on detailed Tencent Cloud Observability Platform alarm policy, please refer to [Create Alarm Policy](#).

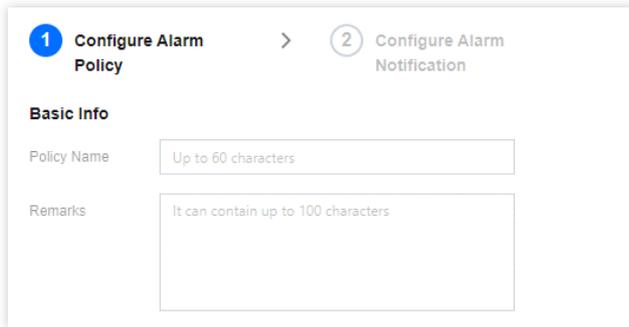
Best Practice

On the cluster details page, select **Cluster Monitoring** on the left to view various performance metrics for the cluster. Click **Set alarms** to jump directly to the following page.



Configuring Basic Information

According to Tencent Cloud Observability Platform rule, FE alarms and BE alarms need to be configured as two separate alarm policies.



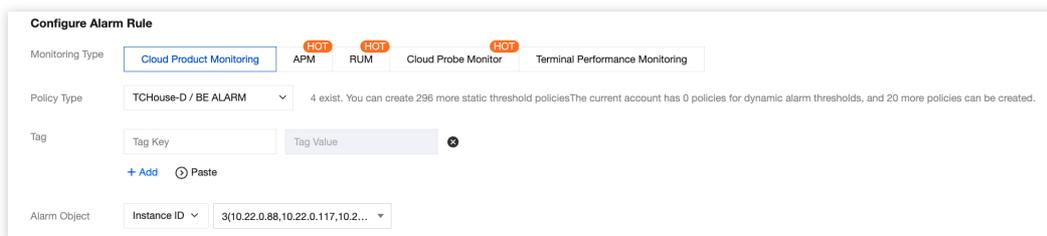
Configuring Alarm Policy

Monitoring Type: Select **Cloud Product Monitoring**.

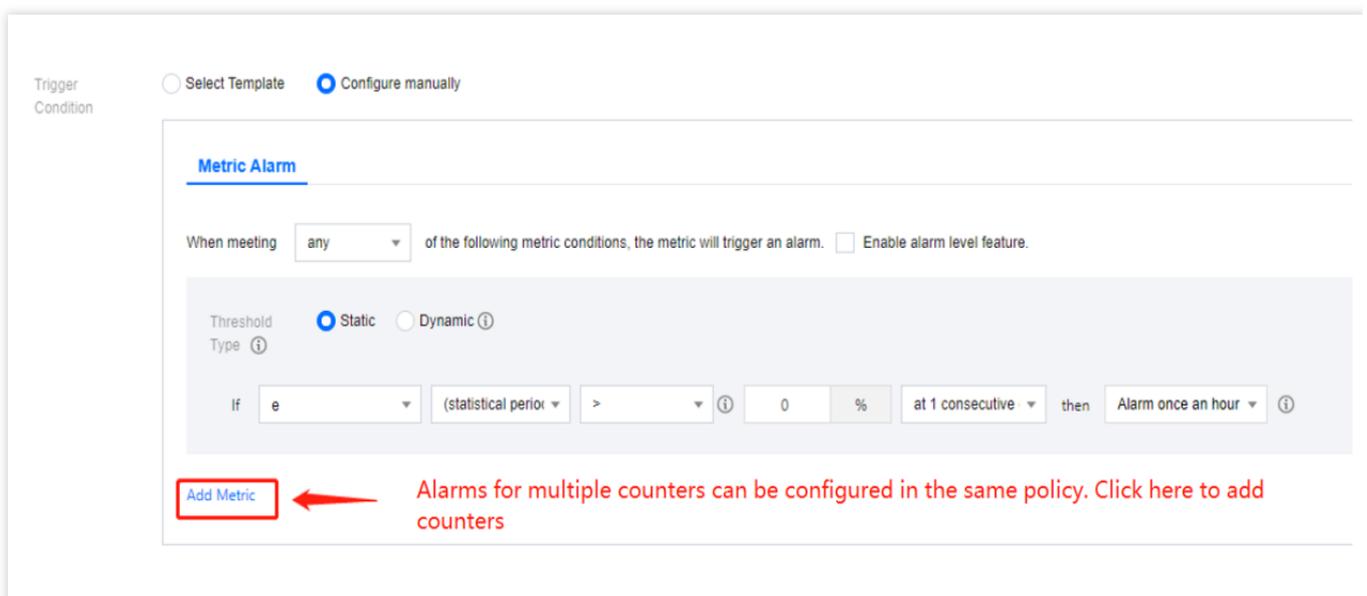
Policy Type: Choose "TCHouse-D/BE ALARM" or "TCHouse-D/FE ALARM" (Note: Alarm for business monitoring metrics needs to be configured under "FE ALARM" type).

Tag: Optional and self-defined.

Alarm Object: Choose the specific alarming node according to the region and instance ID.



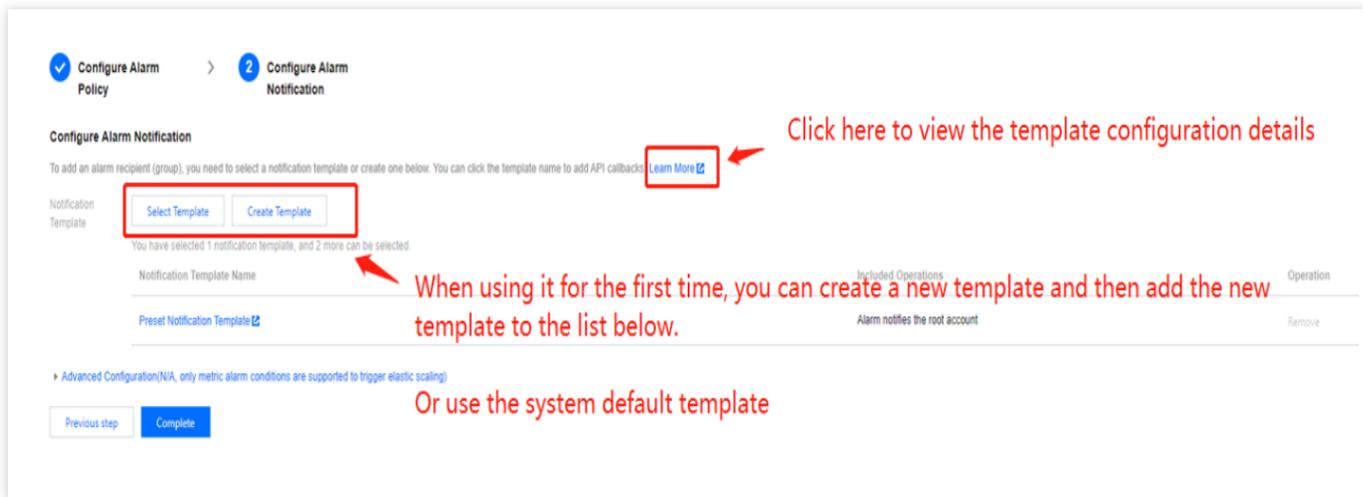
Trigger Condition: Choose to configure manually, and define the rules as needed. Below is an example of configuration:



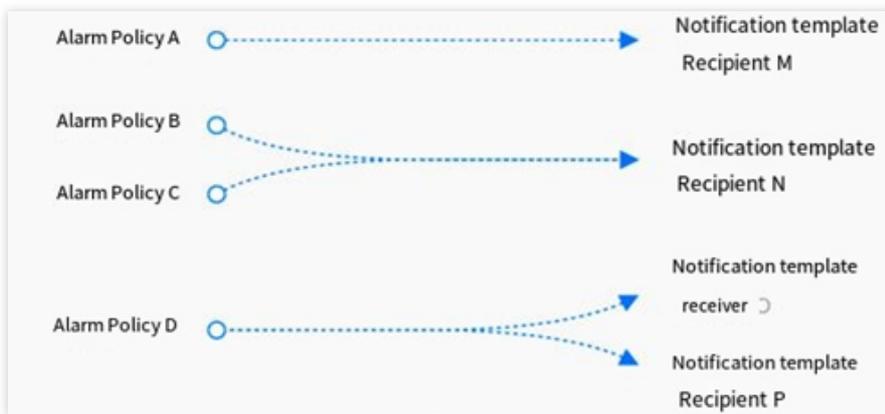
Configuring Alarm Notification

After completing the above alarm rule configuration, you need to specify the scope of alarm notification recipients when an exception occurs. On Tencent Cloud Observability Platform, you need to set the users who are expected to

receive the alarms, API callbacks, etc. as templates, and then associate the templates with the alarms.



On Tencent Cloud Observability Platform, one alarm policy can be bound to multiple notification templates, and one notification template can also be used for multiple alarm policies. You can configure flexibly according to needs.



Viewing Alarm Configuration

After completing the alarm configuration, you can preview existing policies in **Tencent Cloud Observability Platform > Alarm Management > Alarm Configuration** and adjust the alarm policies as needed.

Observability Platform

- Monitor Overview
- Dashboard
- Instance Group
- Report Management
- Alarm Management
 - Alarm Management
 - Alarm Configuration
 - Alarm Governance
 - Alarm Dashboard

Alarm configuration

Alarm Policy NEW Convergence rule Trigger Condition Template Notification Template Scheduling Management

Create Policy Delete More
Advanced Filter
Search by Tag, Policy Name ID

Policy Name	Monitoring Type	Policy Type	Alarm Rule	Project	Associated Instances	Notification Template	Last Modified	Alarm On-Off	Operation
	Tencent Cloud services	TCHouse-O_FE_ALARM	e > 0%, statistical period:1 minute. If the condition is me... cpu_usage > 0%, statistical period:1 minute. If the condi...	-	1	Preset Notification Template	200023294239 2024/11/01 21:39:37	<input checked="" type="checkbox"/>	Copy Delete Alarm f Set to Default Policy
	Tencent Cloud services	TCHouse-O_BE_ALARM	cpu_usage > 0%, statistical period:1 minute. If the condi...	-	4	Preset Notification Template	200023294239 2024/08/01 16:01:52	<input checked="" type="checkbox"/>	Copy Delete Alarm f Set to Default Policy
	Tencent Cloud services	TCHouse-O_BE_ALARM	cpu_usage > 0%, statistical period:1 minute. If the condi...	-	4	Preset Notification Template	200023294239 2024/08/01 15:47:02	<input checked="" type="checkbox"/>	Copy Delete Alarm f Set to Default Policy
	Tencent Cloud services	TCHouse-O_BE_ALARM	e > 0%, statistical period:1 minute. If the condition is me...	-	1	Preset Notification Template	200023294239 2024/03/20 16:00:15	<input checked="" type="checkbox"/>	Copy Delete Alarm f Set to Default Policy

Scan QR code to join our community on WeC

Account Privilege Management

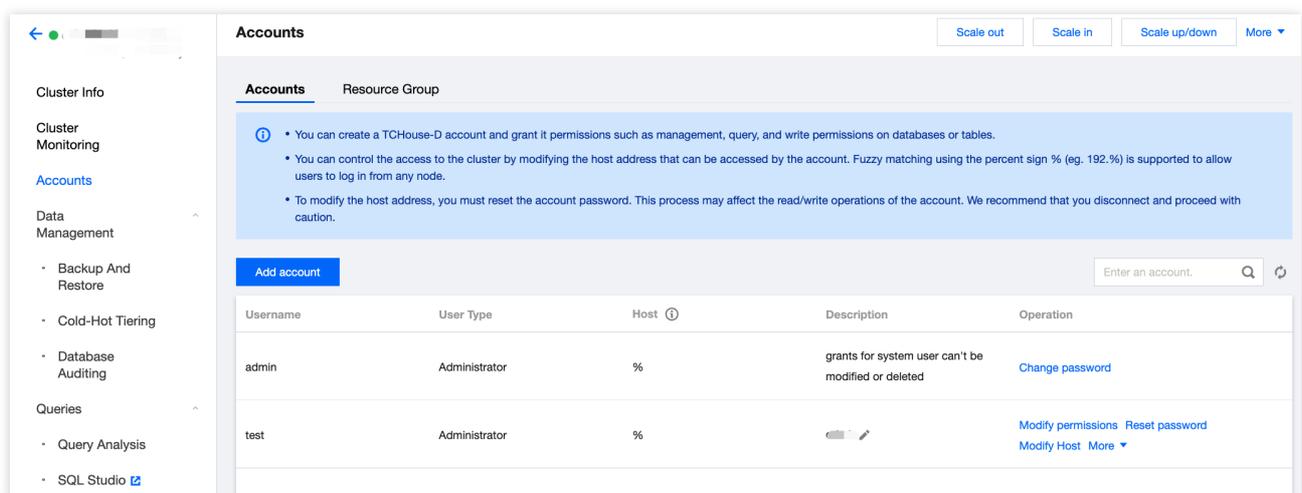
Account Privilege Management

Last updated : 2025-01-16 17:56:09

The Tencent Cloud TCHouse-D Console provides a visual interface for convenient and efficient account and privilege management of clusters.

Account Management

1. Log in to [Tencent Cloud TCHouse-D Console](#), click the target **Cluster ID/Name** and you can find the **Accounts** menu in the left list.
2. The Accounts page provides features for adding/deleting accounts, modifying privileges, resetting passwords.



Adding Account

1. Click the **Add Account** button, fill in the account name, password, host, and description (optional), and click **OK** to add the account.
2. By default, new users have only read privileges for the information_schema library and its tables.
3. User type: You can select Administrator or Regular User. The administrator has access to all databases and tables without additional authorization.
4. Host: Supports individual IP addresses, or use % to imply no restrictions.

Add account ×

Account *

Password * 

Confirm password * 

User Type * ▼

Host ⓘ *

Description

Deleting Account

Deletion is irreversible. Please confirm that this account will not be used in the future before deleting.

After deletion, even if you immediately add an account with the same name, the privileges will be initialized.

Resetting Password

Tencent Cloud TCHouse-D does not allow viewing existing account password, only allows password resets through the console.

If you forget your password, we suggest the following actions:

If you forget the Admin account password, please submit a ticket to [Contact Us](#) to reset the password.

If you forget the sub-account password, you can reset it through the console.

Permission Management

In the user list, click **Modify Permissions** to enter the permission modification window and grant internal table and MultiCatalog permissions to regular users. This window can also be used to **view permissions**.

The administrator has access to all databases and tables by default, without additional authorization.

Internal Table Permissions - Batch Settings for Database/Table Permissions:

You can choose to set internal level permissions in batch: Once the internal level permissions are authorized, the user will have access to all databases and tables under internal.

You can choose to set database-level permissions in batch: Authorize multiple databases at the same time.

You can choose to set table-level permissions in batch: Authorize multiple tables at the same time.

Internal Table Permissions - Column-Level Permissions Setting:

After selecting a specific table, you can grant additional Select permissions for specified fields in this table.

Note: Table-level permissions and column-level permissions will **overlay**. For example, if User 1 does not have permission for Table A, you can grant Select permission for specific columns of Table A to the user.

Muticatalog External Table Permissions:

Modify permissions ✕

CAM User ID: Username:

Host: User Type:

Data Permission Setting

Internal Table Permissions Multicatalog Appearance Permissions

Batch setting for library/table permissions Column-Level permission setting

Data Directory Name	General permissions ?	High-risk permissions ⚠
▼ internal	<input type="checkbox"/> Query <input type="checkbox"/> Insert	<input type="checkbox"/> Alter <input type="checkbox"/> Delete <input type="checkbox"/> Create
Database/Table	General permissions ?	High-risk permissions ⚠
▶ __internal_schema	<input type="checkbox"/> Query <input type="checkbox"/> Insert	<input type="checkbox"/> Alter <input type="checkbox"/> Delete <input type="checkbox"/> Create
▶ aaabbbccddddeefffggghh0	<input type="checkbox"/> Query <input type="checkbox"/> Insert	<input type="checkbox"/> Alter <input type="checkbox"/> Delete <input type="checkbox"/> Create
▶ aaabbbccddddeefffggghh1	<input type="checkbox"/> Query <input type="checkbox"/> Insert	<input type="checkbox"/> Alter <input type="checkbox"/> Delete <input type="checkbox"/> Create
▶ aaabbbccddddeefffggghh10	<input type="checkbox"/> Query <input type="checkbox"/> Insert	<input type="checkbox"/> Alter <input type="checkbox"/> Delete <input type="checkbox"/> Create

Authorization

Authorization scope: Supports granting cluster management privileges, or global, data catalog, database/table privileges.

Granting cluster management privileges:

Once this privilege is activated, it grants the user cluster management privileges (Admin_priv), including query, insertion, modification, deletion, and creation within global scope.

Granting global, data catalog, database/table privileges:

Granting global privileges: Authorization will be effective globally.

Granting data catalog privileges: Authorization will apply to all databases and tables under the data catalog.

Granting database/table privileges: You can specify databases and tables for authorization.

Note:

For external data sources with normal connectivity, we support granting query privileges to the corresponding catalog.

Privilege Category

For internal data sources (internal), the privileges available include standard and high-risk privileges, as categorized below:

Standard privileges:

Query: Read-only privilege (Select) for databases and tables.

Insertion: Write privilege (Load, Insert, Delete) for databases and tables.

High-risk privileges:

Modify: Privilege to alter database tables, including renaming databases/tables, adding/deleting/changing columns, adding/deleting partitions, etc. (Alter).

Delete: Privilege to delete databases, tables, and views (Drop).

Create: Privilege to create databases, tables, and views (Create).

For external data sources (multi-catalog), only query privileges are supported.

Modifying Access Host Address

You can use the console's settings to modify the host address authorized for your account, thereby restricting access to the cluster and enhancing its security.

Note:

The admin account does not support host address modifications.

Setting Host

When creating an account, you must set the host (default is %), support for individual IP address formats, and the use of % implies no restrictions.

Allows for the creation of accounts with the same name but different hosts; privileges, and passwords are independent of each other.

Note:

Supports % for fuzzy matching (e.g., "192.%"), where "%" allows the user to sign in from any node.

Modifying Host

Modifying the host address requires resetting the account's password. The account's read and write access may be affected during the modification process, so it is advised to disconnect with caution.

Modify Host ✕

Resource ID/Name

Account **test**

Host %

New Host ⓘ *

New password * 

Confirm password * 

Privilege Management Through a Client

Last updated : 2024-06-27 10:50:44

This document describes how Tencent Cloud TCHouse-D implements its privilege management from the kernel level.

Overview

The kernel privilege management system refers to MySQL's mechanism, implementing table-level fine-grained control with role-based Access Control, and supports the allowlist mechanism.

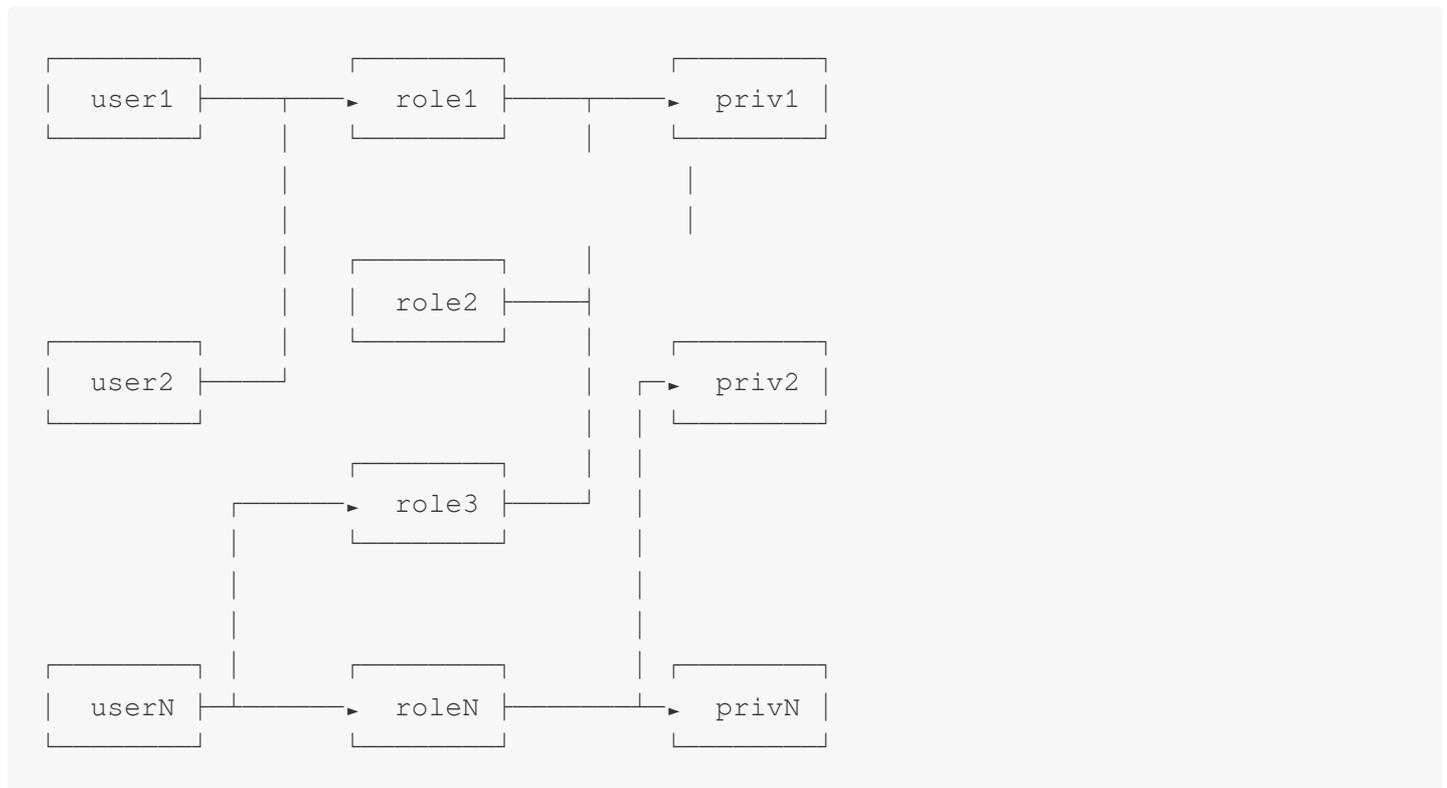
Users and roles are associated, roles and privileges are associated; users are indirectly linked with privileges through roles.

When a role is deleted, users automatically lose all privileges of that role.

When the user and role dissociate, the user automatically loses all privileges of the role.

When the privileges of a role are added or deleted, the privileges of the user also change accordingly.

Authorizing a user is actually authorizing the user's default role. The default role cannot be deleted or assigned to others. When the user is deleted, the default role is also automatically deleted.



Glossary

User_identity

In the privilege system, a user is recognized as a User Identity. User Identity consists of two parts: username and userhost. The username is made up of uppercase and lowercase English characters. Userhost indicates the IP from which the user link comes. User_identity is presented in the form of username@'userhost', indicating username from userhost.

Another representation of user_identity is username@[domain], in which domain is a domain name that can be resolved through DNS or BNS (Baidu Naming Service) into a set of IP. Ultimately they are represented as a group of username@'userhost', so we will use username@'userhost' uniformly.

Privilege

The objects of privileges are nodes, data directories, databases, or tables, and different privileges represent different operation privileges.

Role

Doris can create self-defined named roles. A role can be considered as a collection of privileges. Newly created users can be assigned a role and are automatically granted the privileges that the role has. Subsequent changes in the role's privileges will also be reflected in the privileges of all users belonging to that role.

User_Property

User properties directly subordinate to a user, not to a user identifier. That is, cmy@'192.%' and cmy@[domain] both own the same set of user properties, which belong to user cmy, not to cmy@'192.%' or cmy@[domain].

User properties include but are not limited to: the maximum number of user connections, importing cluster configuration, etc.

Supported Actions

1. Create User: [CREATE USER](#)
2. Modify User: [ALTER USER](#)
3. Delete User: [DROP USER](#)
4. Authorize/Assign Role: [GRANT](#)
5. Revoking Authority/Role: [REVOKE](#)
6. Create Role: [CREATE ROLE](#)
7. Delete Role: [DROP ROLE](#)
8. View Current User Privileges and Role: [SHOW GRANTS](#)
9. View All User Privileges and Roles: [SHOW ALL GRANTS](#)
10. View Created Roles: [SHOW ROLES](#)
11. Set User Properties: [SET PROPERTY](#)
12. View User Properties: [SHOW PROPERTY](#)
13. Modify Password: [SET PASSWORD](#)

For detailed help on the above commands, you can access help in the MySQL client using help + command, such as

```
HELP CREATE USER .
```

Specific Privileges

Privilege Type

Doris currently supports the following privileges:

1. Node_priv

Change privileges for nodes, including the addition, deletion, and offline operations of FE, BE, and Broker nodes.

Root users have this privilege by default. The user who has both Grant_priv and Node_priv can grant this privilege to other users.

This privilege can only be granted at the global level.

2. Grant_priv

Privilege change privilege. Allows to execute operations including authorization, revocation, addition/deletion/change of user/ role etc.

However, users with this privilege cannot grant other users the node_priv privilege unless they themselves have the node_priv privilege.

3. Select_priv

Read-only privileges to databases and tables.

4. Load_priv

Write privileges to databases and tables, including Load, Insert, Delete, etc.

5. Alter_priv

Change privileges to databases and tables, including renaming databases/tables, add/delete/change columns, add/delete partitions, etc.

6. Create_priv

Privilege to create databases, tables, views.

7. Drop_priv

Privilege to delete databases, tables, views.

8. Usage_priv

Privilege to use resources.

Privilege Hierarchy

According to the different scope of privileges, the privileges of the table can be divided into the following four levels:

1. GLOBAL LEVEL: the privileges on `*.*.*` granted by the GRANT statement. The granted privileges are applicable to any table in any database.

2. CATALOG LEVEL: the privileges on `ctl.*.*` granted by the GRANT statement. The granted privileges are applicable to any table in the designated Catalog.

3. DATABASE LEVEL: the privileges on `ctl.db.*` granted by the GRANT statement. The granted privileges are applicable to any table in the designated Database.

4. TABLE LEVEL: the privileges on `ctl.db.tbl` granted by the GRANT statement. The granted privileges are applicable to the designated table in a designated database.

Privilege levels for resources are divided into two levels:

1. GLOBAL LEVEL: the privileges on `*` granted by the GRANT statement. The granted privileges are applicable to resources.

2. RESOURCE LEVEL: the privileges on `resource_name` granted by the GRANT statement. The granted privileges are applicable to a specific resource.

ADMIN/GRANT Privilege Description

The ADMIN_PRIV and GRANT_PRIV privileges owning the **Grant Privilege** are rather special. This document explains each operation related to these two privileges.

1. CREATE USER

The user with ADMIN privileges or GRANT privileges at the GLOBAL and DATABASE levels can create a new user.

2. DROP USER

A user with ADMIN privileges or GRANT privileges at the global level can delete a user.

3. CREATE/DROP ROLE

A user with ADMIN privileges or GRANT privileges at the global level can create a role.

4. GRANT/REVOKE

The user with ADMIN privileges or GLOBAL-level GRANT privileges can grant or revoke the privileges of any user.

The user with CATALOG-level GRANT privileges can grant or revoke the privileges of any user for the specified CATALOG.

The user with DATABASE-level GRANT privileges can grant or revoke the privileges of any user for the specified database.

The user with TABLE-level GRANT privileges can grant or revoke the privileges of any user for the specified table in the specified database.

5. SET PASSWORD

The user with ADMIN privileges or GLOBAL-level GRANT privileges can change the password of any user.

Normal users can change the password for their corresponding UserIdentity. The corresponding UserIdentity can be viewed through the `SELECT CURRENT_USER();` command.

The user with non-GLOBAL-level GRANT privileges cannot change the password of existing users, it can only be specified when creating a user.

Other Descriptions

1. During Doris initialization, the following users and roles are automatically created:

Operator role: This role has Node_priv and Admin_priv, that is, all privileges on Doris.

Admin role: This role has Admin_priv, that is, all privileges except node changes.

root@'%': root user, allowed to log in from any node as operator.

admin@'%': admin user, allowed to log in from any node as admin.

2. It is not supported to delete or change the privileges of the default created role or user.

3. There is only one user with the operator role, that is, Root. Multiple users can be created for the admin role.

4. Description of operations that may cause conflicts:

Domain and IP conflict:

Assuming the following user is created: `CREATE USER cmy@['domain'];`, and authorization: `GRANT SELECT_PRIV ON *.* TO cmy@['domain']`

The domain is resolved to two IPs: ip1 and ip2. Then, we separately authorized cmy@'ip1' once: `GRANT ALTER_PRIV ON *.* TO cmy@'ip1';`, then the privileges of cmy@'ip1' will be modified to SELECT_PRIV, ALTER_PRIV. And when we change the privileges of cmy@['domain'] again, cmy@'ip1' will not change accordingly.

Repeated IP conflict:

Assuming the following users are created: `CREATE USER cmy@'%' IDENTIFIED BY "12345";`, `CREATE USER cmy@'192.%' IDENTIFIED BY "abcde";`

In terms of priority, '192.%' takes precedence over '%', so when the user cmy tries to log in to Doris using the password '12345' from the machine 192.168.1.1, Doris will be rejected.

5. Forget password

If you forget the password and cannot log in to Doris, you can use the following command to log in to Doris without a password on the machine where the Doris FE node is located: `mysql-client -h 127.0.0.1 -P query_port -uroot`, after logging in, you can reset the password through the SET PASSWORD command.

6. No user can reset the password of the root user, except for the root user.

7. The ADMIN_PRIV privilege can only be granted or revoked at the GLOBAL level.

8. In possession of the GLOBAL level GRANT_PRIV is actually equivalent to having ADMIN_PRIV, as the GRANT_PRIV at this level allows the granting of any privilege. Please be cautious.

9. `current_user()` and `user()`

Users can `SELECT current_user();` and `SELECT user();` respectively view `current_user` and `user`. Where `current_user` indicates which identity the current user has been authenticated by the system with, and `user` is the user's actual `user_identity`.

Example:

Assume a `user1@'192.%'` is created, and a user1 from 192.168.10.1 Log in to the system, then the `current_user` is `user1@'192.%'`, while `user` is `user1@'192.168.10.1'`. All privileges are granted to a `current_user`, and the actual user owns all the privileges of the corresponding `current_user`.

Typical Use Cases

Here are some typical use cases of the Doris privileges system.

Scenario 1

Users of the Doris cluster are divided into administrators (Admin), development engineers (RD), and users (Client). The administrator owns all the privileges of the whole cluster, mainly responsible for cluster building, node management etc. Development engineers are responsible for business modeling, including database creation, data import and alteration, etc. Users access different databases and tables to access data.

In this scenario, ADMIN or GRANT privileges can be granted to administrators. RD can be granted CREATE,DROP,ALTER,LOAD,SELECT privileges on any or specific databases and tables. For clients, SELECT privileges can be granted on any or specific databases and tables. Also, different roles can be created to simplify the authorization process for multiple users.

Scenario 2

There are multiple businesses within a cluster, each of which may use one or more data. Each business needs to manage its own users. In such a scenario, the admin user can create a user with the DATABASE level GRANT privileges for each database. This user can only authorize the specified database to users.

Blocklist

Doris itself does not support blocklist, only has an allowlist feature, but we can simulate a blocklist in some ways.

Suppose a user named `user@'192.%'` is created first, indicating users from `192.*` are allowed to sign in. If now we want to prohibit users from `192.168.10.1` from signing in, we can create another user `cmy@'192.168.10.1'`, and set a new password. As the priority of `192.168.10.1` is higher than that of `192.%'`, users from `192.168.10.1` will no longer be able to use the old password to log in.

More help

For more detailed syntax on privilege management, refer to the [GRANTS](#) command manual. You can also input

```
HELP GRANTS
```

 in the MySQL client command line to access more help information.

Data Management

Data Backup

Last updated : 2024-06-27 10:50:59

Tencent Cloud TCHouse-D supports the feature of transferring current data as a file through Broker Backup to a remote storage system and restoring the data to any cluster from the remote storage system. This feature can be used for one-time or regular snapshot backups and also implementations of data migration between different clusters.

Operation Guide

1. Enter [Cluster List](#) and click **Cluster ID/Name** you want to back up.
2. On the cluster management page, click **Data Management > Backup And Restore** on the left to go to the Backup Recovery page. If the Backup service has not been turned on, authorization should be performed first, and a COS bucket for storing data should be selected.

Configure backup service ×

1 **Authorize role** > 2 **Select COS bucket**

i You can configure only one COS bucket for backup.

COS bucket ↕ ↻

Select an existing COS bucket or [create a COS bucket](#)

Confirm **Cancel**

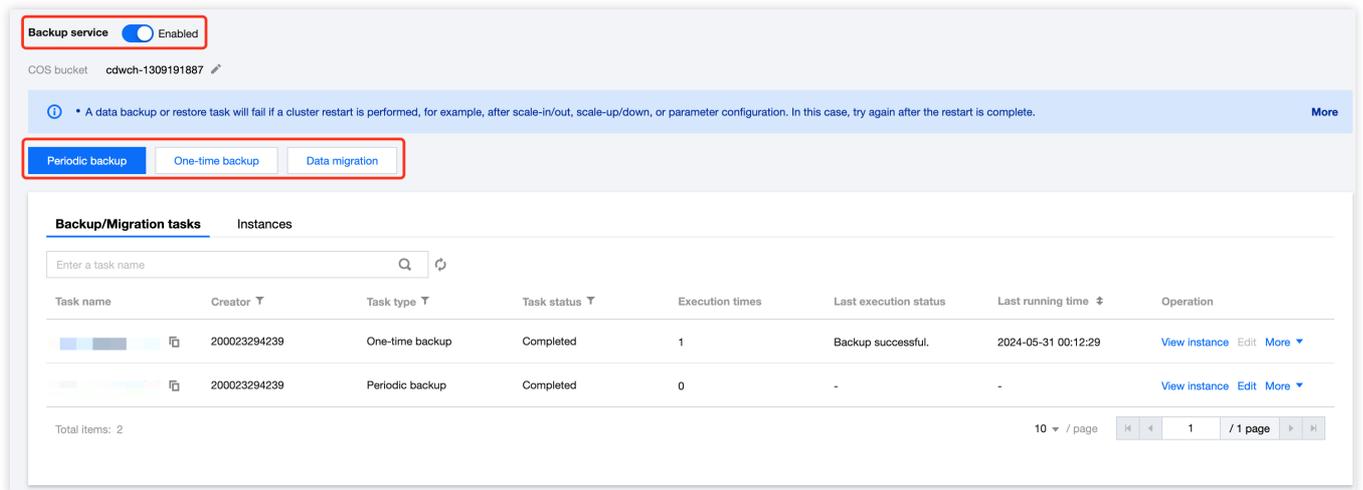
3. After enabling the Backup service, the available operations include: periodic backup, one-time backup, cross-cluster data migration. The Backup feature allows you to choose any database/table to backup, and the data will be stored permanently after a successful backup.

Note:

One-time Backup does not affect the periodic backup policy, but only one BACKUP or RESTORE task can be executed for the same database table at the same time.

The COS bucket path supports changes. After the changes, the new backup snapshot will be written to the new COS bucket, and the backup snapshot stored in the old COS bucket can still be restored.

After backing up the data to the COS, it will be stored as a single copy.



Data Backup Feature Description

You can perform data backup and data migration operations on the Backup Recovery page in Tencent Cloud TCHouse-D. Using the data backup feature, you can perform a one-time backup of the data or plan periodic backup tasks.

Note:

If the Backup or Restore task fails due to cluster restart (such as Scale-out, Scale-in, vertical configuration, parameter configuration operations, etc.), try again after system restart.

One-time Data Backup

When creating a one-time backup task, you can choose the execution method: scheduled execution, or immediate execution.

The granularity of backup can be chosen: backup by table, backup by database, or full backup.

Create one-time backup task

• Only one backup task can be executed for the same database or table at a time. If there is already an ongoing backup task, subsequent tasks will report errors.
• After data is backed up to COS, it will be stored as a single replica.

Scheduling configurations

Task name
6-36 characters; supports Chinese characters, letters, digits, -, and _

Destination
COS bucket

Execution method Scheduled Immediate

Execution period

Backup scope

Granularity Table Database Full

Select Database/Table

Search by keyword

Database/Table	Size
<input type="checkbox"/> __internal_schema	
<input type="checkbox"/> doris_audit_db__	
<input type="checkbox"/> information_schema	
<input type="checkbox"/> mysql	

Selected: 0

Database/Table	Backup size
Select data tables	

Periodic Data Backup

When creating a periodic backup task, you can configure items including: effective cycle, scheduling cycle, and execution time. After the successful creation, the backup task will be executed at the specified time according to the scheduling cycle, within the effective cycle.

One-time Backup does not affect the periodic backup policy, but only one BACKUP or RESTORE task can be executed for the same database table at the same time.

Supports the creation of multiple periodic backup tasks, but if multiple tasks perform backup on the same table at the same time, subsequent tasks will all report errors. In order to ensure data integrity, you can re-backup through the "One-time Backup Data" feature after the previous backup task ends.

Create periodic backup task

Info

- Only one backup task can be executed for the same database or table at a time. If there is already an ongoing backup task, subsequent tasks will report errors.
- After data is backed up to COS, it will be stored as a single replica.

Scheduling configurations

Task name *
6-36 characters; supports Chinese characters, letters, digits, -, and _

Destination
COS bucket

Validity *

Frequency *

Execution period *

Backup scope

Granularity * Table Database Full

Select Database/Table Selected: 0

Database/Table	Size
<input type="checkbox"/> __internal_schema	
<input type="checkbox"/> doris_audit_db_	
<input type="checkbox"/> information_schema	
<input type="checkbox"/> mvsol	

Database/Table	Backup size
Select data tables	

Confirm Cancel

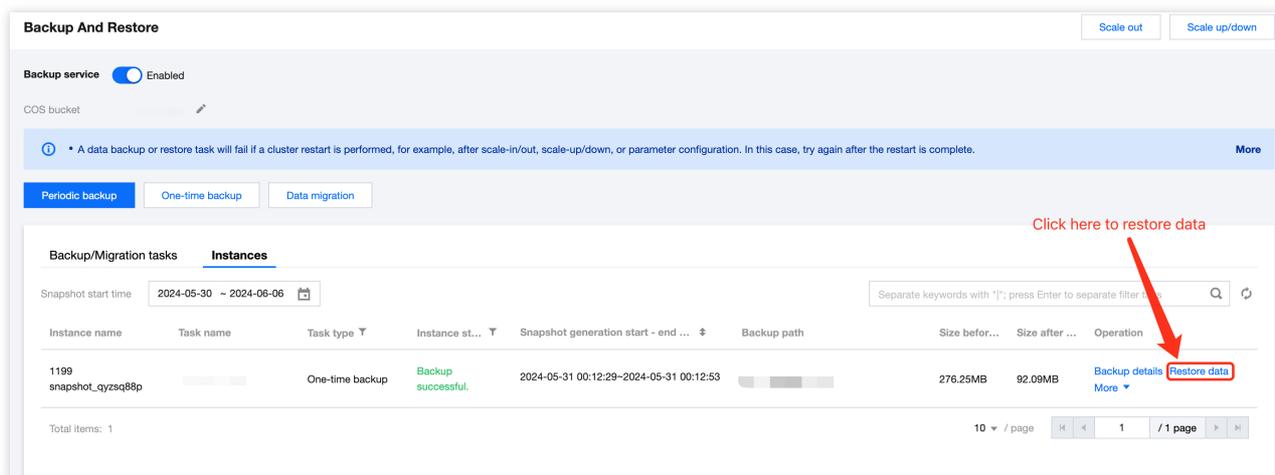
Data Recovery

Last updated : 2024-06-27 10:51:14

Tencent Cloud TCHouse-D supports backing up current data in the form of a file via Broker to a remote storage system, and can restore data to any cluster from the remote storage system. This feature can be used to take one-time or regular snapshot backups of data, or to migrate data between different clusters.

Operation Guide

1. After entering the Backup Recovery page, you can view all tasks in the list below, click **Operation > View Instance** to view all instances under the corresponding task. For instances that have successfully backed up, you can click **Restore data** in **Operation** to perform recovery actions. In **Instances**, you can also view all instances that were recently executed.



Click the Backup Path to view the COS path, COS bucket, snapshot path, and support one-click copy.

2. Click **Restore data** to specifically select the data tables planned for recovery, you can choose whether you want to use the backup table's configuration, including the number of table replicas, dynamic partition open status, etc.

Note:

Backup data will be restored to a database with the same name in the current cluster according to the table dimension, without the need to pre-create a database or table.

During data recovery, if a "table with the same name" exists under the target database, the recovery task will fail. It is recommended to avoid this by changing/deleting the target table.

Restore data ✕

- Backup data will be restored to the current database in the current cluster by table dimension according to the original name of the backup table.
- If there is a "table with the same name" in the data recovery process, the recovery task will fail. If there is a table with the same name, it is recommended to avoid it by changing/deleting the target table.
- After data recovery, the backup table's configurations, including the number of table replicas and the status of the dynamic partition, will be retained.

Restore Scope

Restore Scope Restore by Table Restore by Database Full Restore

Select Database/Table Selected: 0

Search by keyword 🔍

Database/Table	Size
▶ <input type="checkbox"/> doris_audit_db_...	

Database/Table Size

Select data tables

↔

Source Kernel Version 2.0

Target Kernel Version 2.0.7

After recovery, the backup table's entire configurations, including the number of replicas and the status of the dynamic partition, will be retained.

Confirm
Cancel

3. After making your selection, click **Confirm** to start the backup recovery process.

Backup recovery generally requires a certain amount of time to wait, which is related to the amount of data, cluster scale, number of tables, number of shards, degree of parallelism, etc.

To increase the recovery speed, you can contact us to change the configuration through [Ticket](#), but it will require more memory.

Cross-Version Recovery Description

Tencent Cloud TCHouse-D supports restoring data backed up under an older kernel version to a new kernel version cluster, but does not support restoring data backed up in a new version to an older version. The current supported version correspondence is as follows:

Source Cluster Kernel Version	Target Cluster Kernel Version
0.15	0.15, 1.0.X
1.0.X	1.0.X, 1.1.0, 1.1.1, 1.1.2
1.1.0, 1.1.1, 1.1.2	1.1.X
1.1.3, 1.1.4, 1.1.5	1.1.3, 1.1.4, 1.1.5, 1.2.X
1.2.X	1.2.X

Cross-Cluster Data Migration

Last updated : 2024-06-27 10:51:28

Tencent Cloud TCHouse-D supports backing up current data in the form of a file via Broker to a remote storage system, and can restore data to any cluster from the remote storage system. This feature can be used to take one-time or regular snapshot backups of data, or to migrate data between different clusters.

Directions

1. Enter the Backup Recovery page, click **Data Migration**, and fill in the specific configuration information for cross-cluster migration.



2. The source of the data migration can be a Doris cluster or a COS bucket. Fill in the corresponding configuration information. After passing the connectivity check, you can choose the execution method, execution time, and migration range.

If you plan to migrate data from the Doris cluster, you need to fill in information including the source cluster FE_Host, FE_TCP_PORT, and the USER ID and password with Admin privileges.

If you plan to migrate data from a COS bucket, you need to fill in the SecretId, SecretKey of the corresponding COS bucket, and the COS path for the data to be migrated.

Note:

If you want to migrate data from a created Tencent Cloud TCHouse-D cluster and the target and the far end are in the same VPC, it is recommended to use **Doris cluster** as data source.

If you want to cross VPC migration, or migrate backup data from any specified COS bucket, it is recommended to use **COS** as data source.

创建迁移任务

• 迁移会将源端Doris集群/COS桶中的数据、元数据迁移至当前集群。
• 迁移过程中存在“同名表”，迁移任务会失败。如存在同名表，建议通过更改/删除目标端表的方式规避。
• 从COS中迁移数据时，若数据来源为自建Doris集群，请确保备份仓库下快照间不存在库/表重复。

调度配置

迁移任务名称 • 20240204-11-19-迁移
长度限制为6-36个字符，只允许包含中文、字母、数字、-、_

迁移数据来源 Doris集群 COS

源端集群链接

HOST • 请输入DORIS FE HOST	TCP PORT • 9030
USER • 请输入USER	PASSWORD • 请输入PASSWORD

中转COS桶 默认cos桶 terry-test-1301087413 自定义cos桶 ①

执行方式 • 定时执行 立即执行

执行时间 • 选择日期 请选择执行时间

迁移范围

迁移粒度 • 按表迁移 按库迁移 全量迁移

选择数据库/表 已选择 0 张表

- After completing the configuration of the data source for the migration, you can set the execution method and time, and choose the library, table to be migrated, and click **Migrate** to execute.
- When you are migrating data from the Doris cluster, by default, the data will be synchronized to the COS bucket on the target end. You can choose to restore data after the backup is completed during the configuration.
- When you are migrating data from a COS bucket, the data will be automatically synchronized to the Doris cluster at the target end after the migration is completed.

Note:

Backup data will be restored to the database with the same name in the current cluster by table dimension, and there is no need to create a library table in advance.

Cross-Cluster Migration Feature Description

Migrating from Doris Cluster

If you want to migrate data between two clusters in the same VPC, you can set the Doris cluster as the source of the data migration. To establish a connection with the remote cluster, you need to fill in information including the source cluster FE_Host, FE_TCP_PORT, and the USER ID and password with Admin privileges.

After the migration is completed, the data of the remote Doris cluster will be synchronized to the COS bucket designated by the target cluster. You can also configure your own defined COS bucket, specify the migration COS path, and migrate the data to the CosPath/doris_manage_backup path.

迁移数据来源 Doris集群 COS

源端集群链接

HOST * 请输入DORIS FE HOST TCP PORT * 9030

USER * 请输入USER PASSWORD * 请输入PASSWORD

连接

中转COS桶 默认cos桶 terry-test-1301087413 自定义cos桶 ⓘ

SecretId * 请输入SecretId SecretKey * 请输入SecretKey

请输入CosPath 连接检查

You can choose whether to restore the data after migration completion. If checked, the data will be migrated to the specified COS bucket and also automatically synchronized to the target Doris cluster.

Migrating from COS Bucket

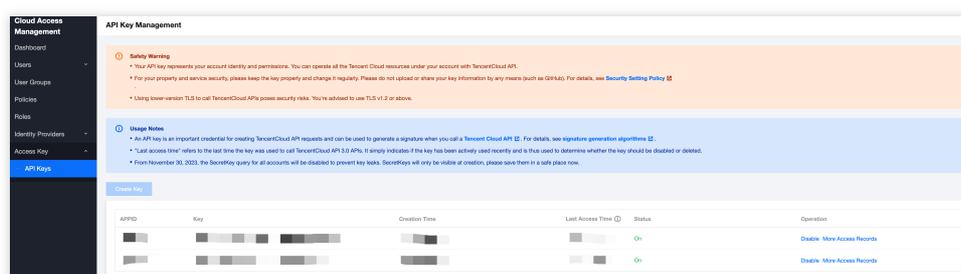
Once you have completed the data backup in the remote cluster, the backup instances will be stored in the selected COS bucket, and you can implement cross-cluster data migration based on this.

Note:

The data migration feature does not support cross-region. It is suggested to copy the data in the remote COS bucket to the COS bucket in the target region before migration.

To initiate a backup task that migrates from COS, you need to choose COS as the source of the migration data, and fill in the corresponding key of the remote COS bucket and the COS path of the data to be migrated.

User-defined key: You can create and manage user-defined keys on the [API Key Management](#) page. This key can be used to manage all the Tencent Cloud resources under your account. Keep it safe.



The COS path to be migrated is formatted as `https://$domain/$backup_dir/$snapshot_dir`.

`$domain`: domain name (can be found in the Overview page of the corresponding bucket in [Cloud Object Storage](#)).

`$backup_dir`: The path of the folder where the file is located.

`$snapshot_dir`: Backup version snapshot folder path.

After the migration, the data will be automatically synchronized to the target Doris cluster.

Cross-Version Migration Version Description

Tencent Cloud Warehouse-D supports migrating data from old kernel version clusters to new kernel version clusters, but does not support migrating data backed up from new versions to old versions. The supported version correspondence is as shown in the following table:

Source Cluster Kernel Version	Target Cluster Kernel Version
0.15	0.15, 1.0.X
1.0.X	1.0.X, 1.1.0, 1.1.1, 1.1.2
1.1.0, 1.1.1, 1.1.2	1.1.X
1.1.3, 1.1.4, 1.1.5	1.1.3, 1.1.4, 1.1.5, 1.2.X
1.2.X	1.2.X

Database Auditing

Last updated : 2024-12-20 15:49:51

The Database Auditing feature displays the historical information of all SQL operations within the cluster and supports search, filtering, and download.

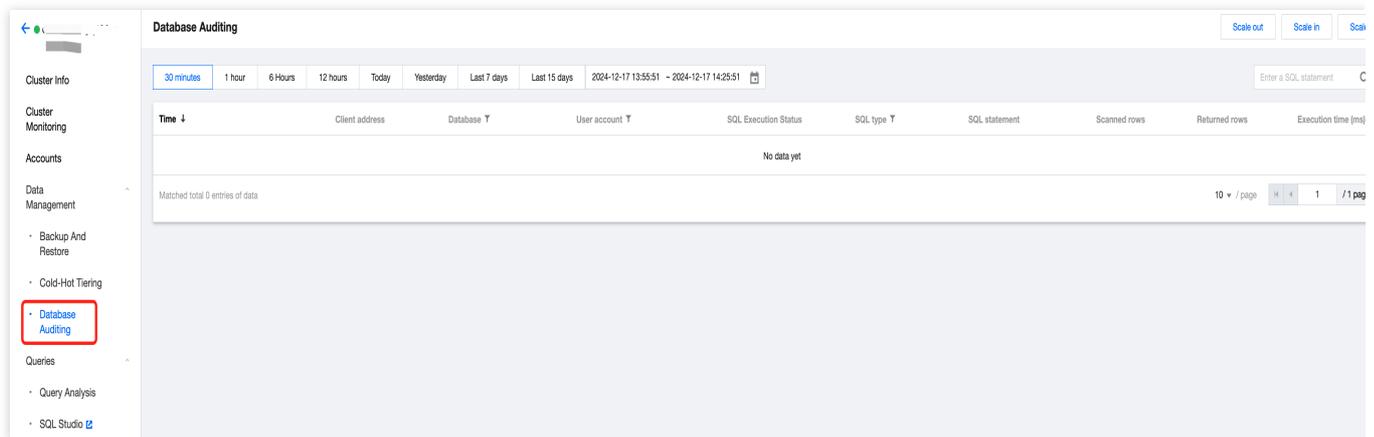
Information displayed includes: time, client address, database name, user account, SQL type, SQL statement, returned rows, scanned rows, and execution time.

Note:

To ensure data security, Tencent Cloud Warehouse-D currently does not support collecting audit logs via ES/CLS. If needed, you can query the audit tables `doris_audit_tbl__` in the default library `doris_audit_db__` for analysis.

Operation Guide

Log in to [Tencent Cloud TCHouse-D Console](#), and select **Data Management > Database Auditing** from the left side list to enter the audit page.



Select a time range: Supports customizing the time range to seek any time frame's audit log.

Search and filter: To precisely locate query records, you can filter the audit logs by the SQL statement. You can also filter by database name, user account, SQL type, and more.

Download data: Support exporting audit logs within a certain time range, with a single export limit of 20,000 records. If this limit is exceeded, adjust the time range.

Download database audit log ✕

Cluster ID

Log time 

 Up to 20,000 log entries can be exported at a time. To export more entries, please set an appropriate time range and try again.

Note:

During download, the conditions set by the user for search and the filters for database name, user account, and SQL type columns on the page will be retained.

Cold-Hot Data Tiering

Last updated : 2025-01-16 17:56:09

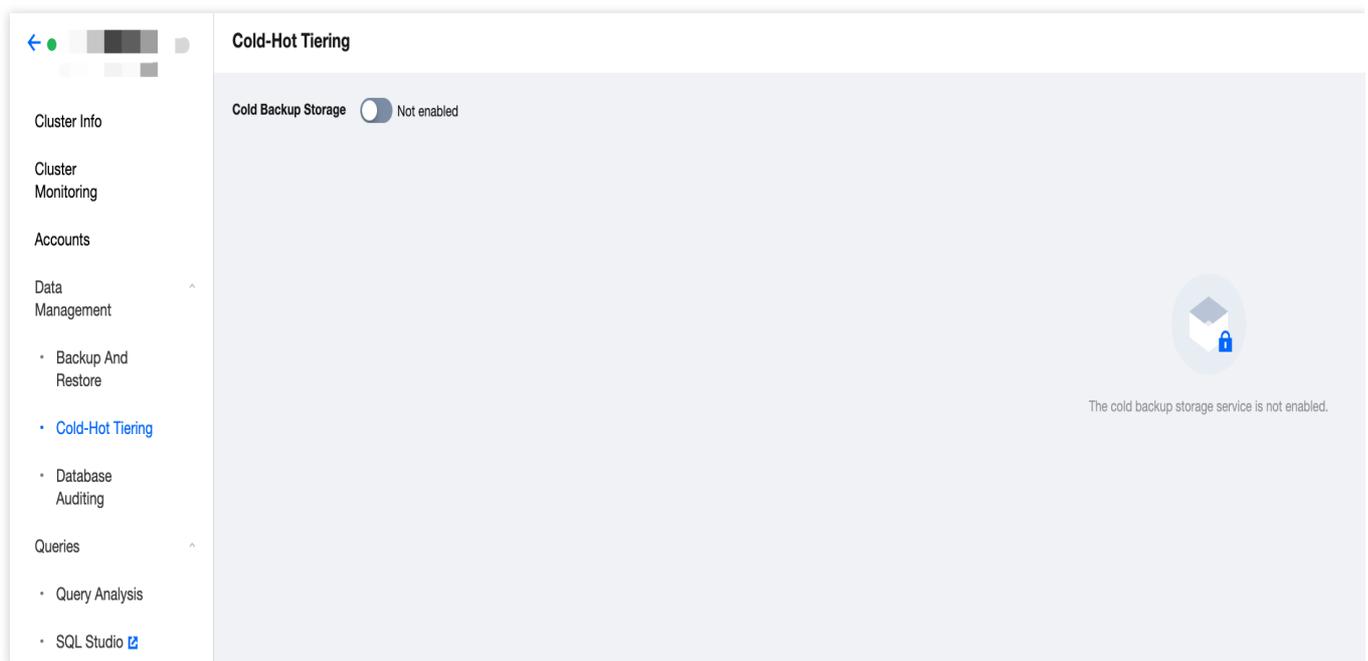
Tencent Cloud TCHouse-D supports data cooling with lower frequency of inquiry to COS, greatly reducing the storage cost of cold data.

Note:

Currently, only clusters with 2.0/2.1 kernels support this feature. This feature requires allowlist usage. If necessary, [submit a ticket](#) to contact us.

Operation Guide

1. Go to [Cluster List](#) and click the **Cluster ID/Name** of the cluster where you want to configure the cold-hot data tiering policy.
2. On the Cluster Management page, click **Data Management > Cold-Hot Data Tiering** on the left to enter the Cold-Hot Data Tiering page and click **Switch** to enable cold-hot data tiering.



3. When you perform data cooling, the cold data will be written to COS, and therefore you need to select a COS bucket for data storage.

Note:

Before enabling the cold-hot data tiering feature, check the lifecycle policy of the COS bucket to avoid impacts on query performance. The Cloud Storage Object will automatically convert storage types or delete objects based on the rules you set within the specified time frame.

Once the cold-hot data tiering feature is enabled, it cannot be disabled, and the COS storage address cannot be modified.

Enable Cold Backup Storage ✕

✓ Authorize role > **2 Select COS bucket**

i Only COS buckets in the same region as the cluster are available for hot-cold tiered storage due to network restrictions.

COS bucket

Select an existing COS bucket or [create a COS bucket](#)

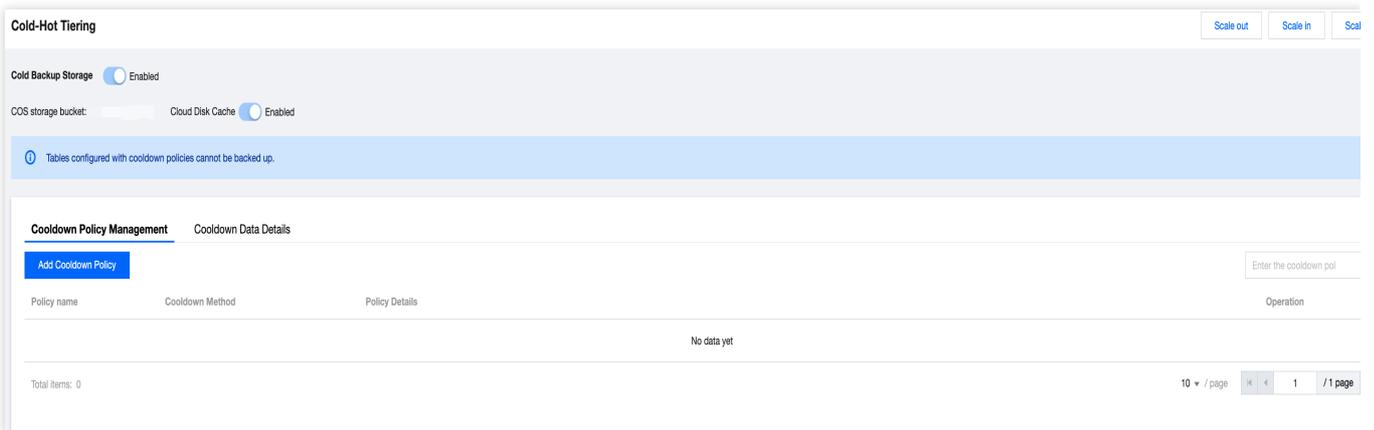
When the cold-hot tiering feature is enabled, to avoid impacts on query performance, please check the lifecycle policy of the COS bucket. The Cloud Storage Object will automatically convert storage types or delete objects based on the rules you set within the specified time frame.

Confirm **Cancel**

4. After the cold-hot data tiering feature is successfully enabled, you can manage the cooling policies and view the cooling details on the page.

Note:

To improve the performance of querying cold data, the cloud disk cache is enabled by default. The larger the cloud disk cache, the better the query performance. If you need to adjust the cache size, [submit a ticket](#) to contact us.



5. Click **Create Cooling Policy** button to configure data cooling policies. The following two cooling methods are supported:

Performing cooling operations on outdated data that has not been updated for a long time: Once this policy is associated, outdated data (at the segment file level) that has not been updated beyond **cooling TTL** will be converted to cold data and stored in COS, while new data will continue to be written to the BE cloud disk.

Perform overall cooling from a specified time: Once this policy is associated, all tables/partitions associated with this policy will be cooled and stored in COS at the **cooling time**, and subsequent data will also be directly cooled and stored in COS.

Create Cooldown Policy ✕

i Cooldown policy can be edited after configuration, but it cannot be deleted.

Policy name *

Cooldown Method * Cool down old data that has not been updated for a long time.
 Cool down entirely from the specified time

Cooldown TTL * Hour

Policy Details **After this policy is associated, any data not updated within the Cooldown TTL will be converted to cold data and stored in Cloud Object Storage (COS). New data will pursue to be stored in BE disks as hot data.**

6. Select a cooling policy and click **Associate with Databases and Tables** to implement batch association of databases/tables or partitions to the same cooling policy. Note that system tables do not support cooling policies.

Associate Cooldown Policy with Database and Table



- Cooldown policies can be associated by table or partition. If a hot and cold data separation policy is set for a table, different hot and cold data separation policies cannot be set for partitions of the table.
- After a table/partition is associated with a cooldown policy, the association cannot be changed to ensure consistent cold/hot data. Please operate with caution.

Cooldown Granularity * Table-based Cooldown Partition-based Cooldown

Select Database/Table

Search by keyword	
Database/Table	Size
▼ __internal_schema	
audit_log	738.69KB
column_statistics	0B
histogram_statistics	0B
▶ information_schema	
▶ mysql	

Selected: 0

Database/Table
audit_log

Partition Selection

Enter a keyword or regular expression.

- p20241204
- p20241205
- p20241206
- p20241207

7. On the **Cooling Data Details** tab page, you can view the cooling data details of the node, database, and table granularity.

Cold-Hot Tiering Scale out Scale in Scale

Cold Backup Storage Enabled

COS storage bucket: Cloud Disk Cache Enabled

ⓘ Tables configured with cooldown policies cannot be backed up.

Cooldown Policy Management Cooldown Data Details

View Method By Node By Database and Table

BE Node IP	Disk Space	Space Used by Hot Data	Space Used by Cooldown Data
10.22.0.46	199.862 GB	247.040 KB	0.000
10.22.0.102	199.862 GB	231.488 KB	0.000
10.22.0.103	199.862 GB	270.171 KB	0.000

Query Management

Query Analysis

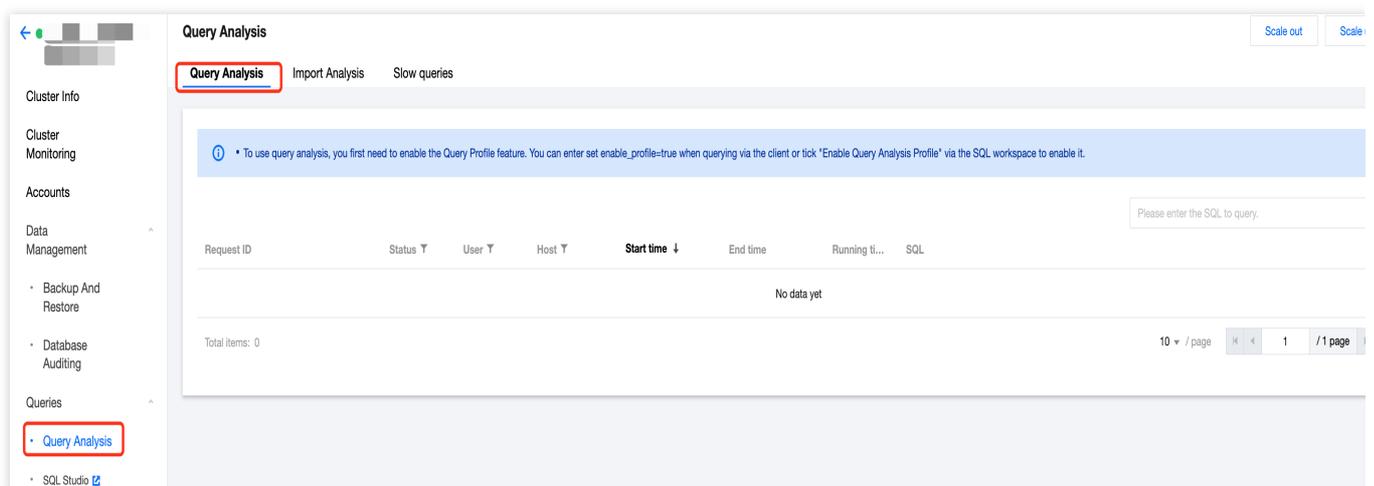
Last updated : 2024-06-27 10:52:03

On the Query Analysis page, you can analyze query details of those with the Profile feature enabled. On this page, you can view SQL details, query plan, and execution details, presented in both graphical and textual format.

With these capabilities, developers can understand query execution in finer detail and carry out targeted debugging and optimization.

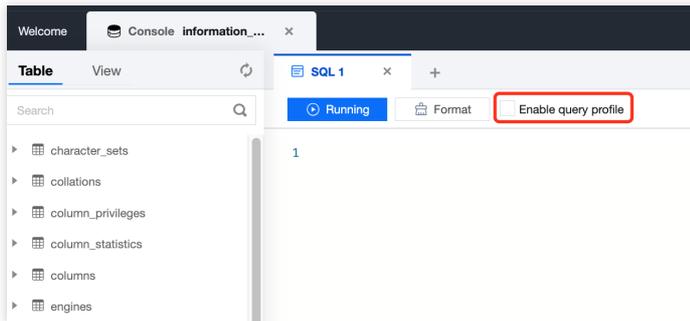
Operation Guide

1. Log in to [Tencent Cloud TCHouse-D Console](#), select the target **Cluster ID/Name** from the cluster list, and switch to the Query Analysis page.



2. The queries that you want to analyze need to have the Profile feature enabled in advance. The enabling methods include SQL workspace enabling and client enabling:

Enable SQL Studio: When using **SQL Studio** for queries, you can check **Enable query profile** while editing the query.

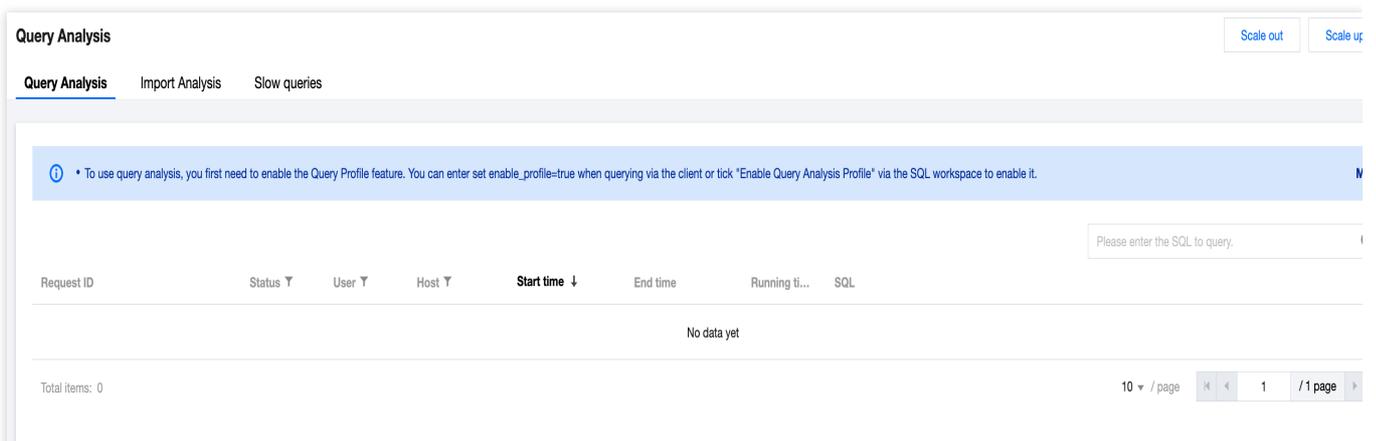


Enable Client: When using other methods for queries, enter set enable_profile=true in the command line to enable the Profile feature.

Note:

Query and Analysis only applies to SQL with the Profile feature enabled. This feature may have some impact on query performance.

3. For profile enabled queries, you can view Request ID, Query Status, Operation Time, etc. in the Query and Analysis list. By clicking **Request ID** or **Details** button in the SQL column, you can proceed with detailed analysis.

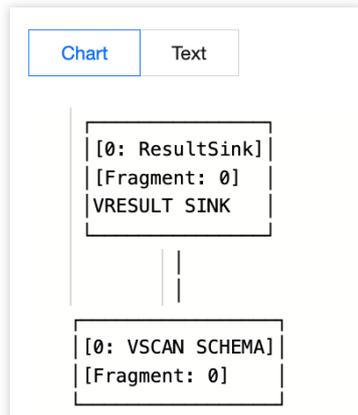


4. On the Query Analysis detail page, you can view into the SQL Details, Query Plan, and Execution Details of the corresponding Request ID:

SQL Details: You can view the specific statement of the executed SQL, with support for copying with a single click.

Query Plan: It presents the current query plan in graphic or text format.

Execution Details: It presents the execution details of a query in graphic or text format.

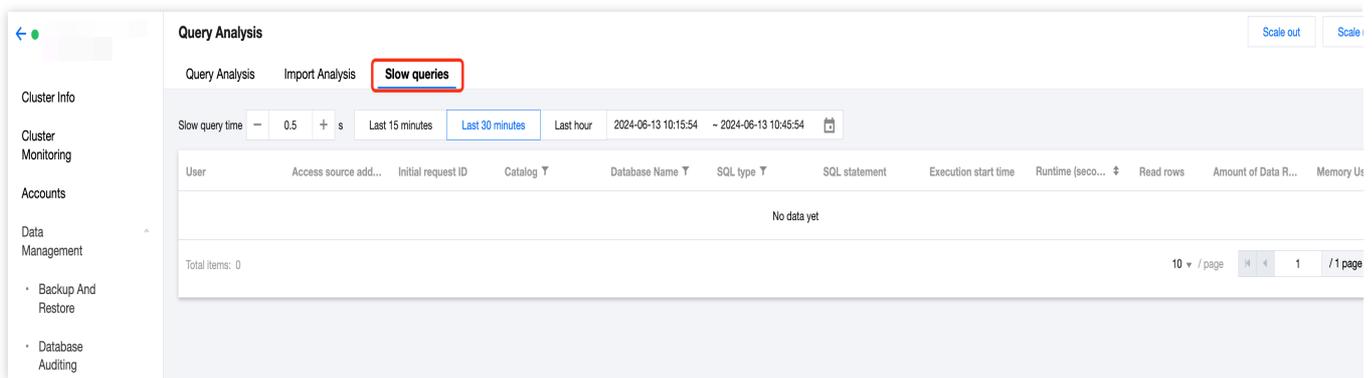


Slow Query Analysis

Last updated : 2024-06-27 10:52:17

Operation Guide

1. Log in to [Tencent Cloud TCHouse-D Console](#), and select the target cluster by **ID/Name** in the cluster list to enter the cluster details page.
2. After switching to the Query Analysis page, you can switch to **Slow Queries** sub-page to perform slow query analysis.



3. This page, by default, displays slow SQL taking more than 500 ms. You can adjust the threshold for judging slow query SQL duration as needed, supporting a minimum setting of 100 ms.

Features

You can set the time range for analysis, including the last 15 minutes, last 30 minutes, last hour, or a custom time period.

It supports sorting the slow queries list by execution duration, or filtering based on database name, or SQL types:

Database Name: If you did not specify the database during the query, this column will be displayed as NULL.

SQL Type: query and non-query.

Query Analysis

Query Analysis Import Analysis **Slow queries**

Slow query time: 0.5 s | Last 15 minutes | **Last 30 minutes** | Last hour | 2024-06-13 10:15:54 ~ 2024-06-13 10:45:54

User	Access source add...	Initial request ID	Catalog	Database Name	SQL type	SQL statement	Execution start time	Runtime (seco... ⚡	Read rows	Amount of Data R...	Memory Usa
Total items: 0											

10 / page | 1 / 1 page

- All
- Non-Query
- Query

OK Reset

Modify Configurations

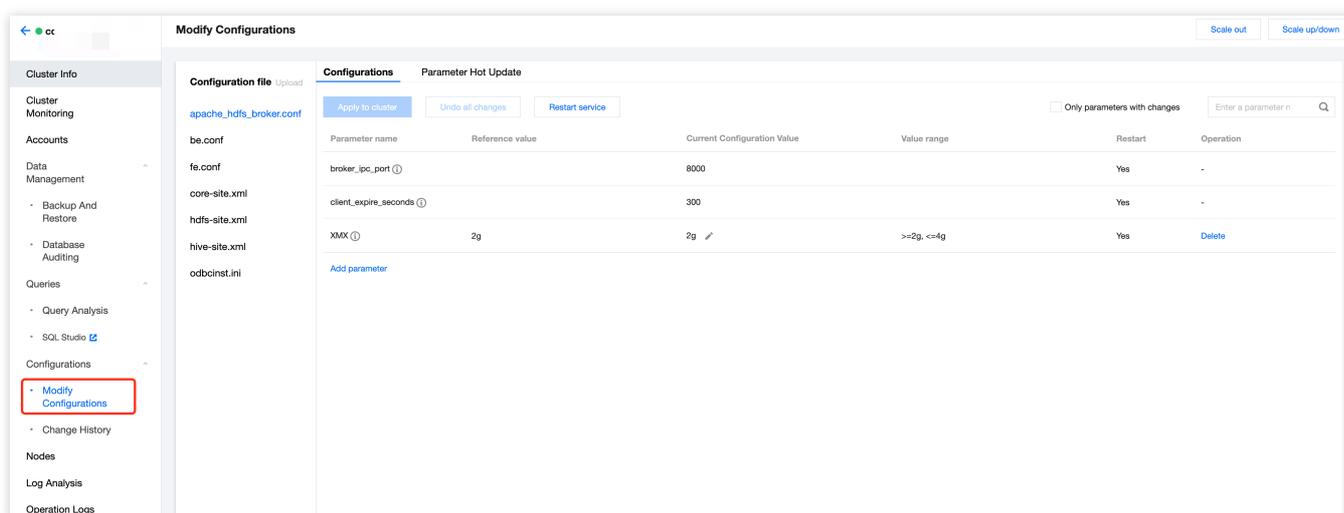
Modify Configurations Through Console

Last updated : 2024-06-27 10:52:32

Tencent Cloud TCHouse-D supports flexible parameter configuration methods. You can comprehensively and intuitively set parameters such as BE/FE/Broker in the console and support to view historical modification records.

Operation Guide

1. Log in to [Tencent Cloud TCHouse-D Console](#), and click **Cluster ID/Name** in the cluster list to view cluster details. Select **Configurations> Modify Configurations** from the left sidebar to enter the parameter configuration page.



2. The configuration parameters corresponding to Broker, BE and FE support configuration management, and parameter hotspot update. Whole file modification is supported for XML configuration files.

Note:

To ensure the stability of the service, the feature to upload configuration files has been disabled. If you need to add new configuration files, contact us by [Submit a ticket](#).

Modifying Configuration Parameters

The configuration parameters corresponding to Broker, BE, and FE support configuration management, and the parameter hot updates. Configuration modification will be applied to the node after a restart, and parameter hot updates will take effect immediately without restarting.

Configuration Management

For parameters in the conf configuration file, you can view the reference value, the current configuration value, and the range of modifications. Parameters in the conf configuration file can be added, deleted, or modified.

Modifying parameters: You can define parameters within a certain range. After editing, click **Apply to Cluster** to save and distribute. The modifications will take effect after rebooting the respective node by clicking **Restart Service**.

Adding parameters: You can add new configuration parameters in the corresponding configuration file. Like modifying parameters, the new parameters will take effect after saving, distributing, and restart.

Deleting parameters: Some parameters support deletion. After saving, distributing, and rebooting the corresponding node, deletion will take effect. The cluster will run according to the default parameters. Be cautious.

The screenshot shows the 'Modify Configurations' interface. At the top, there are buttons for 'Apply to cluster' and 'Restart service'. Below is a table of configurations with columns for Parameter name, Reference value, Current Configuration Value, Value range, Restart, and Operation. A modal is open for editing the 'XMX' parameter, showing a 'Running value' of 12g and a rule of >=8g, <=16g. The modal has 'Confirm' and 'Cancel' buttons.

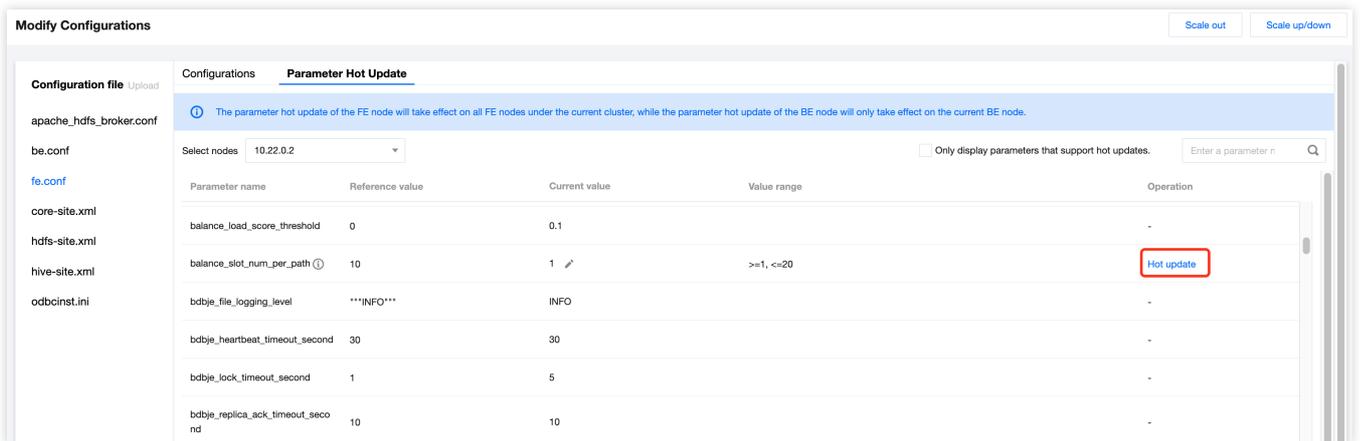
Configuration file	Parameter name	Reference value	Current Configuration Value	Value range	Restart	Operation
be.conf	audit_log_enable_compress	false	true		Yes	-
fe.conf	edit_log_port		9010		Yes	-
core-site.xml	http_port		8030		Yes	-
hdfs-site.xml						
hive-site.xml						
odbcinst.ini						
	JAVA_OPTS_FOR_JDK_9		-Dlog4j2.formatMsgNoLookups=true -Xmx\$XMX -XX:SurvivorRatio=8 -XX:MaxTenuringThreshold=7 -XX:+CMSClassUnloadingEnabled -XX:-CMSParallelRemarkEnabled -XX:CMSInitiatingOccupancyFraction=80 -XX:SoftRefLRUPolicyMSPerMB=0 -Xlog.gc:\$LOG_DIR/fe.gc.log.\$DATE		Yes	-
	lower_case_table_names		0		Yes	-
	query_port		9030		Yes	-
	rpc_port				Yes	-
	sys_log_enable_compress	false			Yes	-
	sys_log_level	INFO			Yes	-
	XMX	12g	12g	>=8g, <=16g	Yes	Delete

Parameter Hot Update

For parameters in the configuration file, you can view the reference values, current running values, and modification ranges. After modifying parameters, click **Hot Update** to apply modification. The parameter hot update will take effect immediately after the modification, with no cluster restart needed.

Note:

The parameter hot update for the FE node will apply to all FE nodes in the cluster, and the parameter hot update for the BE node will only apply to the current BE node.

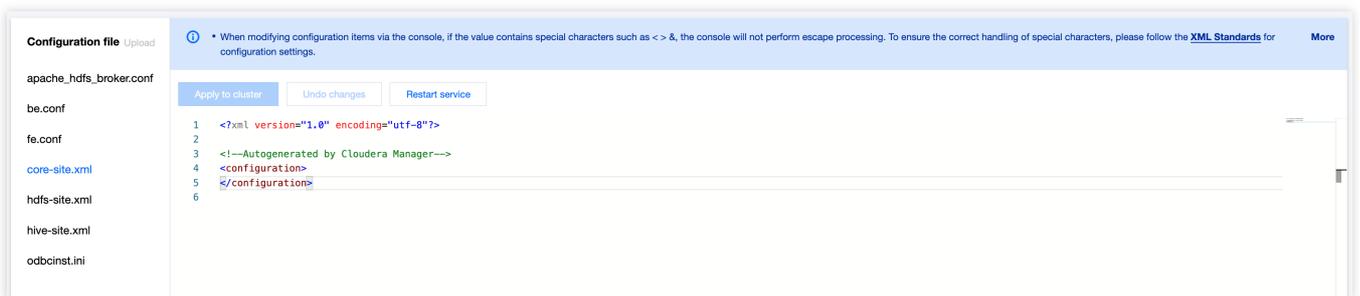


Modifying XML File

For XML configuration files, you can edit the entire file as needed. Click **Apply to Cluster** to save and distribute modification. The configuration modifications will take effect after restarting the service.

Note:

While modifying the configuration items, if the value contains special characters like <, >, & etc., the console won't escape them. To ensure the correct handling of special characters, follow [XML Standards](#) for configuration.



View Modification History

In the cluster list, click **Cluster ID/Name** to view cluster details. Select **Configurations > Change History** from the left sidebar to enter the configuration modification records page.

Supports viewing the parameter modification records within any selected time range. You can specifically view the modification time, the account ID of the person who made the changes, and the change description.

If you want to understand the differences in the configuration file before and after a specific modification, you can click **Operation > Comparison** for detailed viewing.

Change History Scale out Scale up/down

Today | Last 7 days | Last 15 days | Last 30 days | 2024-06-13 ~ 2024-06-13

Configuration file	Modified at	Modified by	Change description	Operation
be.conf	2024-06-13 11:30:41	200023294239	-	Completion

Total items: 1

10 / page 1 / 1 page

- Cluster Info
- Cluster Monitoring
- Accounts
- Data Management
 - Backup And Restore
 - Database Auditing
- Queries
 - Query Analysis
 - SQL Studio
- Configurations
 - Modify Configurations
 - Change History**

Modify Configurations Through Client

Last updated : 2024-06-27 10:52:46

In addition to the Console, you can also configure parameters through the client. This document will introduce related configuration methods.

FE Parameter Configuration

The configuration file of FE, `fe.conf`, is usually placed in the `conf/` directory of the FE deployment path. Another configuration file, `fe_custom.conf`, is used to record configuration items that users dynamically configure and persist at runtime.

After the FE process starts, it will first read the configuration items in `fe.conf`, and then read the configuration items in `fe_custom.conf`. The configuration items in `fe_custom.conf` will overwrite the same configuration items in `fe.conf`.

The location of the `fe_custom.conf` file can be configured through the `custom_config_dir` configuration item in `fe.conf`.

Viewing Configuration Items

After FE starts, you can view the configuration items of FE in the MySQL client using the following command: `ADMIN SHOW FRONTEND CONFIG;` The meanings of each column in the results are as follows:

Key: name of the configuration item.

Value: current value of the configuration item.

Type: value type of the configuration item, whether it's integer or character string.

IsMutable: Whether it can be dynamically configured. If true, it means that this configuration item can be dynamically configured at runtime. If false, it means that this configuration item can only be configured in `fe.conf` and will take effect after restarting FE.

MasterOnly: Whether it's a configuration item unique to the Master FE node. If true, it means that this configuration item only makes sense in the Master FE node and has no meaning for other types of FE nodes. If false, it means that this configuration item is meaningful in all FE nodes.

Comment: description of the configuration item.

You can also open the FE front page `http://fe_host:fe_http_port/variable` in your browser. You can view the currently effective FE configuration items in Configure Info.

Setting Configuration Items

There are two ways to configure FE's configuration items:

1. Static Configuration

Add and set configuration items in the `conf/fe.conf` file. The configuration items in `fe.conf` will be read at the start of the FE process. Configuration items not in `fe.conf` will use the default value.

2. Dynamic Configuration Through the MySQL Protocol

After FE starts, you can set configuration items dynamically with the following command. This command requires administrative rights.

```
ADMIN SET FRONTEND CONFIG ("fe_config_name" = "fe_config_value");
```

Not all configuration items support dynamic configuration. By using the `ADMIN SHOW FRONTEND CONFIG;` command, the `IsMutable` column of command results reveals whether it supports dynamic configuration.

Configuration items modified in this way will be invalidated after the restart of the FE process.

For more help with this command, the `HELP ADMIN SET CONFIG;` command can provide information.

3. Dynamic Configuration Through the HTTP Protocol

For more specific details, please refer to [Set Config Action](#). This method can also persistently modify the configuration items. The configuration items will be persisted in the `fe_custom.conf` file, and they will still be effective after FE restarts.

Application Example

1. Modifying `async_pending_load_task_pool_size`

With `ADMIN SHOW FRONTEND CONFIG;`, you can see that this configuration item cannot be dynamically configured (`IsMutable` is false). You need to add `async_pending_load_task_pool_size=20` in `fe.conf`, then restart the FE process to apply this configuration.

2. Modifying `dynamic_partition_enable`

With `ADMIN SHOW FRONTEND CONFIG;`, you can see that this configuration item can be dynamically configured (`IsMutable` is true) and it is a configuration unique to Master FE. First, we can connect to any FE and execute the following command to modify the configuration:

```
ADMIN SET FRONTEND CONFIG ("dynamic_partition_enable" = "true");`
```

Then you can use the following command to check the modified value:

```
set forward_to_master=true;
ADMIN SHOW FRONTEND CONFIG;
```

After you modify it in this way, if the Master FE restarts or a Master switch happens, the configuration will become invalid. This configuration can be permanently enabled by directly adding the configuration item in `fe.conf` and restarting FE.

3. Modifying `max_distribution_pruner_recursion_depth`

With `ADMIN SHOW FRONTEND CONFIG;`, you can see that this configuration item can be dynamically configured

(`IsMutable` is true) and it is not a configuration unique to Master FE.

Similarly, we can modify this configuration through the command to modify the configuration dynamically. Since this configuration is not unique to the Master FE, you need to connect separately to different FEs and dynamically modify the configuration. This ensures that all FEs use the modified configuration value.

Configuration Item List

For a detailed list of configuration items applicable to FE, see [FE Configuration Items](#).

BE Parameter Configuration

The configuration file for BE, `be.conf`, is usually located in the `conf/` directory of the BE deployment path. Another configuration file, `be_custom.conf`, is used to record configuration items that users dynamically configure and persist during running.

After the BE process starts, it will first read the configuration items in `be.conf`, and then read the configuration items in `be_custom.conf`. The configuration items in `be_custom.conf` will overwrite the same items in `be.conf`.

Viewing Configuration Items

Users can view the current configuration items by accessing the BE Web page at

```
http://be_host:be_webserver_port/varz
```

Setting Configuration Items

There are two ways to configure BE's configuration items:

1. Static Configuration

Add and set configuration items in the `conf/be.conf` file. The configuration items in `be.conf` will be read when BE starts. Items not in `be.conf` will use the default values.

2. Dynamic Configuration

After BE starts, the following command can be used to dynamically set configuration items.

```
curl -X POST http://{be_ip}:{be_http_port}/api/update_config?{key}={value}'
```

The following command can be used to persist modifications to the configuration. The modified configuration items will be stored in the `be_custom.conf` file.

```
curl -X POST http://{be_ip}:{be_http_port}/api/update_config?{key}={value}&persist=true
```

Application Example

1. Modifying `max_base_compaction_concurrency` in a static way

Add `max_base_compaction_concurrency=5` in the `be.conf` file, and then restart the BE process to apply the configuration.

2. Modifying `streaming_load_max_mb` dynamically

After BE starts, use the following command to dynamically set the configuration item `streaming_load_max_mb` :

```
curl -X POST http://{be_ip}:{be_http_port}/api/update_config?
streaming_load_max_mb=1024
```

If the return value is as follows, the setting is successful.

```
{
  "status": "OK",
  "msg": ""
}
```

The configuration will lose its effect after the BE is restarted. If you want to persist modifications, use the following command:

```
curl -X POST http://{be_ip}:{be_http_port}/api/update_config?
streaming_load_max_mb=1024\&persist=true
```

Configuration Item List

For a detailed list of configuration items applicable to BE, see [BE Configuration Items](#).

User Configuration Items

The User level configuration is only valid for individual users. Each user can set their own user properties, which will not affect each other.

Viewing Configuration Items

After FE starts, users can view a User's configuration items in the MySQL client using the following command:

```
SHOW PROPERTY [FOR user] [LIKE key pattern] ; For specific syntax, use the command help show
property; .
```

Setting Configuration Items

After FE starts, in the MySQL client, use the following command to modify a User's configuration items:

```
SET PROPERTY [FOR 'user'] 'key' = 'value' [, 'key' = 'value'] ; For specific syntax, use the
command help set property; .
```

User-level configuration items will only be effective for specified users and will not affect the configuration of other users.

Application Example

1. Modify user Billie's `max_user_connections` .

Use `SHOW PROPERTY FOR 'Billie' LIKE '%max_user_connections%';` to view the current maximum number of connections for user Billie, which is 100.

Use `SET PROPERTY FOR 'Billie' 'max_user_connections' = '200';` to change the current maximum number of connections for user Billie to 200.

Configuration Item List

For a detailed list of user configuration items, see [User Configuration Items](#).

Node Management

Last updated : 2024-06-27 10:53:07

Feature Overview

The Nodes page displays the status list for various roles in the cluster, which include FE, BE, BROKER.

For each role, the health status, process status, node IP and last restart time are displayed (which corresponds to the time of the last manual service restart operation on the Role Management page).

On the Nodes page, you can perform restart, and start operations on each node. The types of restart include immediate restart, graceful restart, and rolling restart. Batch execution for the same type of nodes is supported.

Role	Node health status	Node IP	Last restarted
<input type="checkbox"/> FE(master)	Healthy	10.22.0.2	-
<input type="checkbox"/> BROKER	Healthy	10.22.0.2	-
<input type="checkbox"/> FE(observer)	Healthy	10.22.0.6	-
<input type="checkbox"/> BROKER	Healthy	10.22.0.6	-
<input type="checkbox"/> FE(observer)	Healthy	10.22.0.11	-
<input type="checkbox"/> BROKER	Healthy	10.22.0.11	-
<input type="checkbox"/> BE	Healthy	10.22.0.12	-
<input type="checkbox"/> BROKER	Healthy	10.22.0.12	-
<input type="checkbox"/> BE	Healthy	10.22.0.3	-
<input type="checkbox"/> BROKER	Healthy	10.22.0.3	-
<input type="checkbox"/> BE	Healthy	10.22.0.35	-
<input type="checkbox"/> BROKER	Healthy	10.22.0.35	-

To ensure the high availability of each node, in addition to user's manual operations, each node runs a daemon process. If any node is detected to have an abnormal exit, the daemon process will immediately launch the corresponding node.

Note:

The restart operation generally takes 20-30 seconds. The page will automatically refresh to access the latest status during the execution.

Your operations on the Nodes page will be recorded on the Operation Logs page for any time reference.

Log Analysis

Last updated : 2024-06-27 10:53:32

Tencent Cloud TCHouse-D supports log collection, and you can enable Log Service for clusters as needed. Log Service supports both ES and CLS forms.

Note

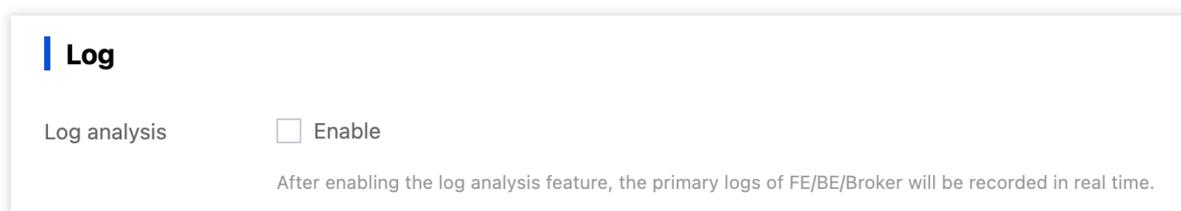
ES service currently only supports the Beijing, Shanghai, Guangzhou. For billing information, see [ES Billing Overview](#). CLS is a type of postpaid service. Please keep sufficient funds to avoid affecting your log upload and display. For billing information, see [Log Service Billing Overview](#).

Enabling Log Service

You can enable Log Service when creating a new cluster or at any time while the cluster is running. Such operations are explained in detail below.

Enabling During New Cluster Creation

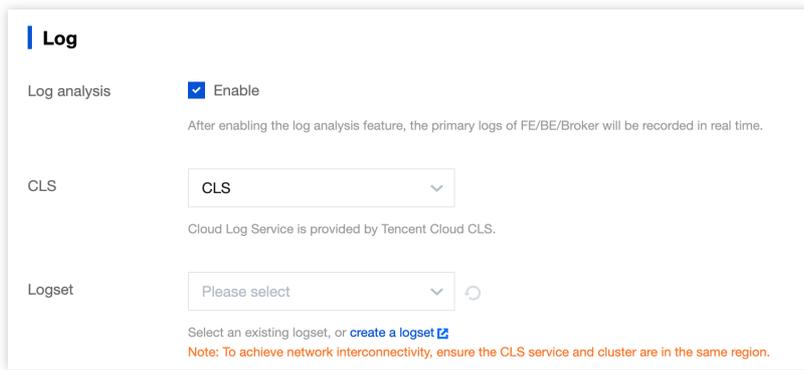
During the creation of a cluster, you can choose to enable Log Service as needed. If you choose ES Log Service (currently supports Beijing, Shanghai, Guangzhou regions), you need to select a project space for management. If you have never created it before, click **Create new project space** to quickly create it. The log analysis feature is provided by ES Serverless service. For details, see [ES Serverless Service Overview](#).



If you choose CLS Log Service, you need to select or create a CLS log set in the same region as the cluster. CLS logs are stored for 30 days by default.

Note

A new log Topic will be created in the log set you configured. You can preview your log Topic on the CLS Log Service page. Do not delete the configured log topic; otherwise it will cause the log retrieval page to query failure.



Log

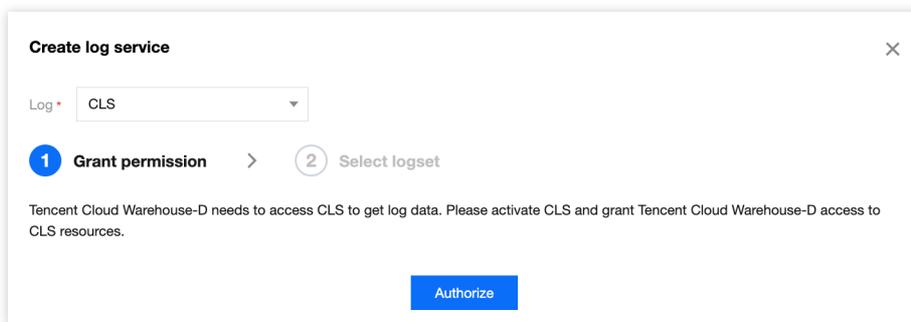
Log analysis Enable
After enabling the log analysis feature, the primary logs of FE/BE/Broker will be recorded in real time.

CLS
Cloud Log Service is provided by Tencent Cloud CLS.

Logset
Select an existing logset, or [create a logset](#).
Note: To achieve network interconnectivity, ensure the CLS service and cluster are in the same region.

Enabling or Modifying Existing Clusters

If you didn't enable Log Service when creating the cluster, you can log in to the [Tencent Cloud TCHouse-D Console](#) at any time and configure the Log Service through **Operation > More > Create log service** configuration in the cluster list. For unauthorized users, you need to authorize first, and then configure the log set.



Create log service

Log

1 Grant permission > **2 Select logset**

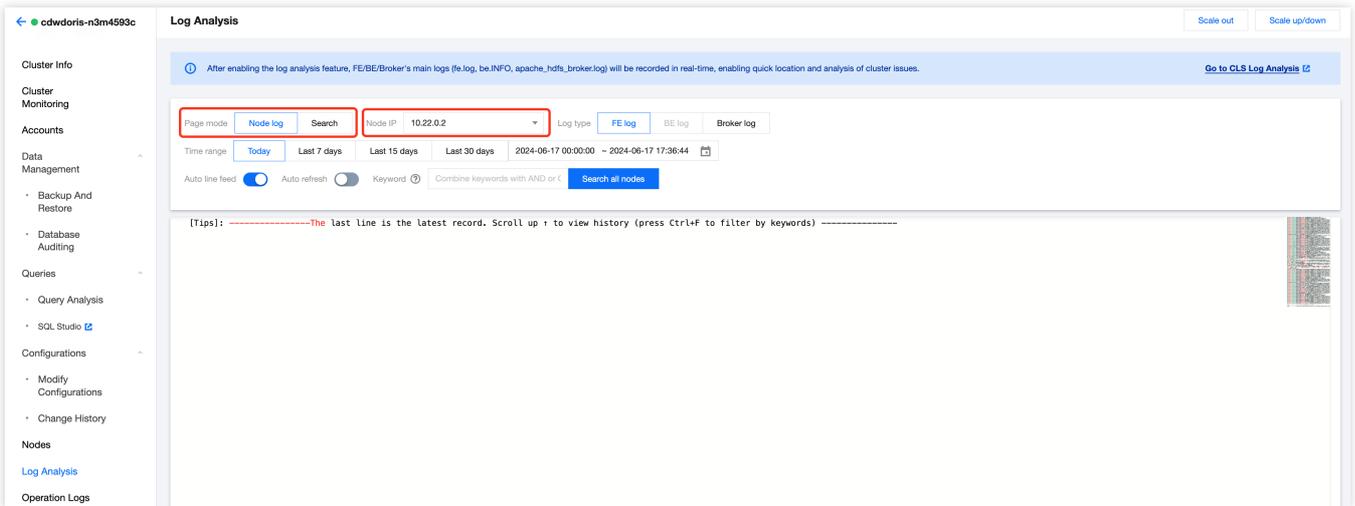
Tencent Cloud Warehouse-D needs to access CLS to get log data. Please activate CLS and grant Tencent Cloud Warehouse-D access to CLS resources.

Log Search

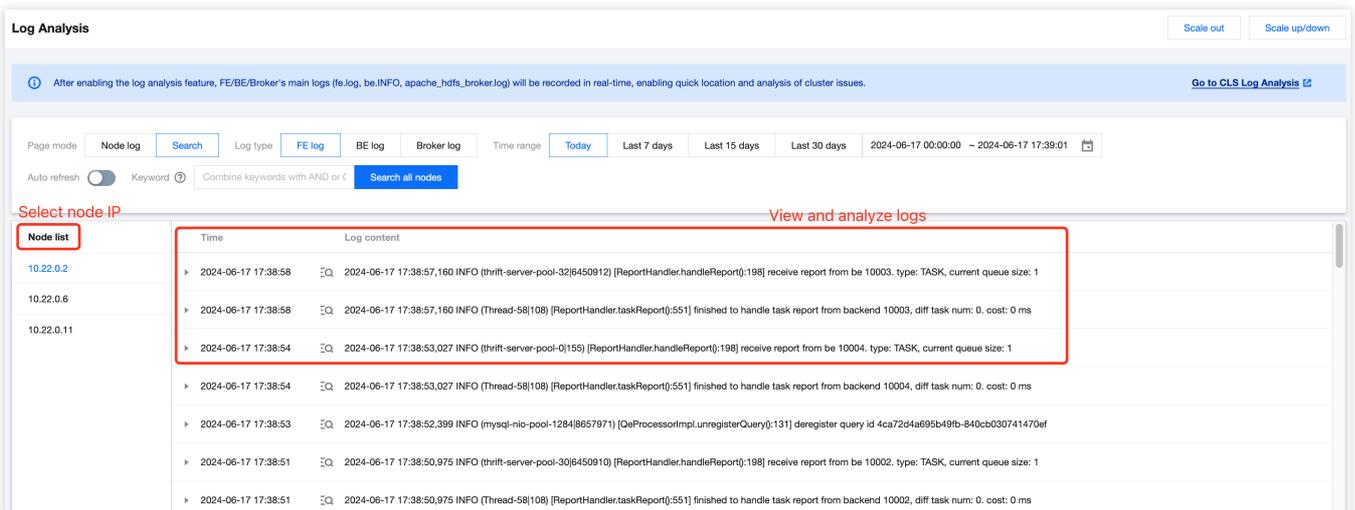
After enabling Log Service, you can easily retrieve logs from the console to assist in operations.

Operation Instruction

1. Log in to [Tencent Cloud TCHouse-D Console](#), click **Cluster ID/Name** from the **Cluster List** to enter cluster details page, and switch to the **Log Analysis** page.
2. Both **Node Log** and **Search** page modes are supported. In the **Node Log** mode, you can view logs by node IP and log type, and perform log analysis using features such as specifying time ranges and keyword search.



3. In **Search mode**, you can conveniently set search keywords, retrieve logs within a selected time range, and group the filtered results by node IP. The syntax and rules of keywords are detailed in [Syntax and Rules](#). The query results display the most recent 100 records for each node by default.



How to Correctly Read FE/BE Log?

In many cases, problems need to be investigated through logs, so here we explain the format and view methods of FE/BE logs.

FE Log

FE log mainly includes:

fe.log: main log. It includes all content except fe.out.

Fe.warn.log: subset of the main log. It records only WARN and ERROR level logs.

A typical example of an FE log is as follows:

```
2021-09-16 23:13:22,502 INFO (tablet scheduler|43)
[BeLoadRebalancer.selectAlternativeTabletsForCluster():85] cluster is balance:
default_cluster with medium: HDD. skip
```

Where:

`2021-09-16 23:13:22,502` : log time.

`INFO` : log level, INFO by default.

`(tablet_scheduler|43)` : thread name and thread id. Through the thread id, we can view the context information of this thread and facilitate troubleshooting.

`BeLoadRebalancer.selectAlternativeTabletsForCluster():85` : class name, method name, and line number.

`cluster is balance xxx` : log content.

BE Log

BE log mainly includes:

be.INFO: main log. This is actually a soft link, linked to the latest be.INFO.xxxx.

be.WARNING: subset of the main log. It only records WARN and FATAL level logs. This is a soft link, linked to the newest be.WARN.xxxx.

A typical example of an BE log is as follows:

```
I0916 23:21:22.038795 28087 task_worker_pool.cpp:1594] finish report TASK.  
master host: 10.10.10.10, port: 9222
```

`I0916 23:21:22.038795` : log level and date and time. Capital letter I indicates INFO, W indicates WARN, and F indicates FATAL.

`28087` : Thread id. With the thread id, you can view the context information of the thread, making it easier to troubleshoot the thread.

`task_worker_pool.cpp:1594` : code file and line number.

`finish report TASK xxx` : log content.

SQL Studio

Last updated : 2024-06-27 10:53:48

Through the SQL Studio, you can quickly connect to the cluster and carry out a series of operations using SQL commands.

Operation Guide

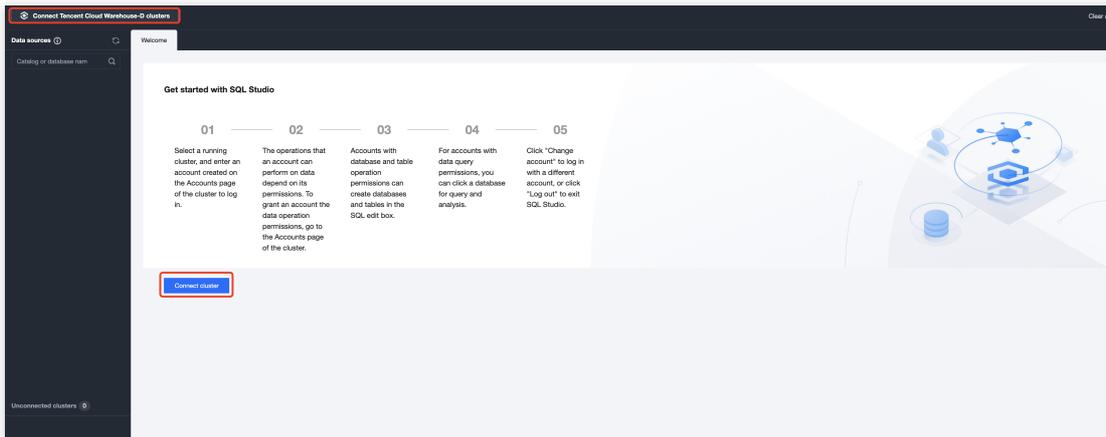
1. After logging in to [Tencent Cloud TCHouse-D Console](#), you can access **SQL Studio** by clicking on the left list. Select the target **Resource ID/Name** in the cluster list. You can also find the entrance under Queries.

The screenshot displays the 'Cluster Info' page in the SQL Studio console. The page is divided into several sections:

- Basic info:** Cluster ID: cdwdoris-lavdpwz4, Cluster name: test12M2, Billing mode: Pay-as-you-go, Creation time: 2024-03-27 10:39:44, Tag: Doris.
- Configuration info:** Kernel version: 1.2, High availability (HA): Read high availability, FE node spec: Standard, 4-core 16 GB / 3 node(s) / CLOUD_SSD 200 GB, BE node spec: Standard, 8-core 32 GB / 5 node(s) / CLOUD_PREMIUM 200 GB, Table names are case-sensitive: Case-Insensitive (Convert all table names to lowercase for storage.), Cloud Block Storage Encryption: Enable, Security Groups: sg-ksnuusa.
- Cluster status:** Cluster status: Serving.
- Network info:** AZ: ap-singapore-2, VPC ID, Subnet ID, JDBC access address, HTTP connection address, and a table of Node info.

No.	Node type	Node health status	Node IP	Cluster VPC IP
1	FE(master)	Healthy	10.0.1.152	9.0.16.42
2	FE(observer)	Healthy	10.0.1.159	9.0.16.15
3	FE(observer)	Healthy	10.0.1.10	9.0.16.16
4	BE	Healthy	10.0.1.219	9.0.16.3
5	BE	Healthy	10.0.1.238	9.0.16.37
6	BE	Healthy	10.0.1.99	9.0.16.26
7	BE	Healthy	10.0.1.147	9.0.16.35
8	BE	Healthy	10.0.1.243	9.0.16.28

2. Once you enter the SQL Studio, you first need to connect to a cluster. Data operations depend on the data privileges of the logged in account, which can be granted in the Accounts of the corresponding cluster. See [Privilege Management Through the Console](#) for details.



3. When connecting to a cluster, you need to select a region, cluster ID, and enter a valid account and password. Click **Accounts** to go to the console page for account operation.

Note:

Please use the account with host % to log in to SQL Studio.

Connect cluster ✕

i • Use an account whose host is % to log in to the SQL workspace.

Region: 🌐 ap-singapore 4 Other regions 6

Cluster: test12M2

Account: admin ↻

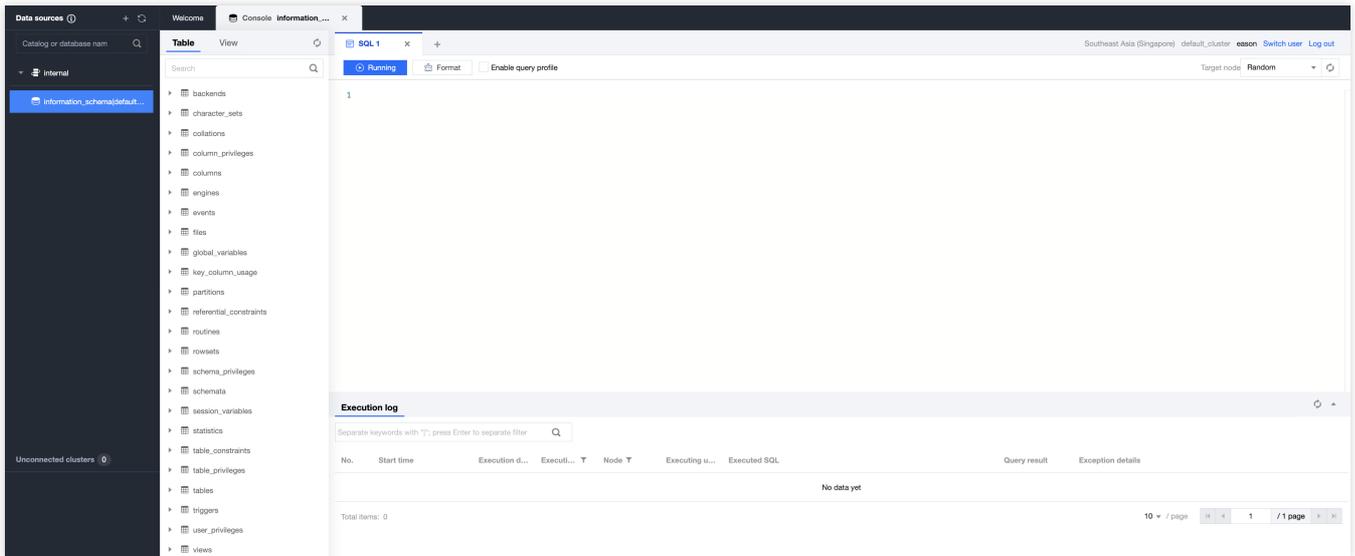
[Accounts](#) 🔗

Password:

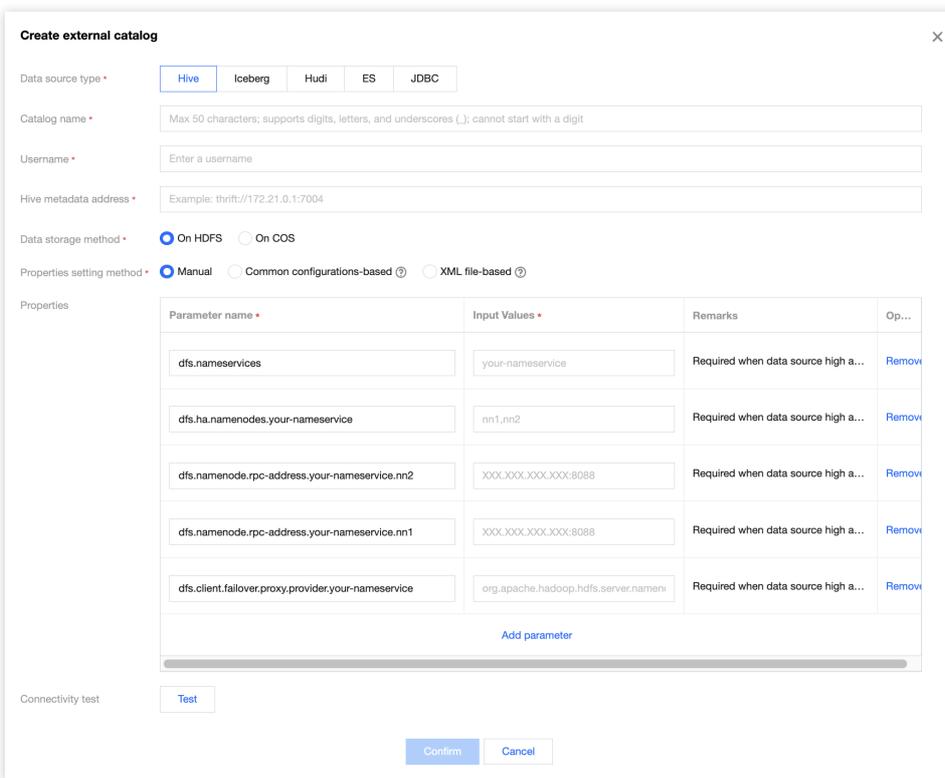
Data connectivity: Test

Confirm
Cancel

4. In SQL Studio, you can view the SQL statement execution records. Search and filtering are supported.



5. Clicking the button in the left list allows you to easily create external data sources.



Note:

The tables of built-in data sources can be created through DDL statements in the SQL compilation box, and external data sources only support queries.

6. Click **Switch User** to switch accounts. After use, click **Log Out** to exit SQL Studio.

Enabling Resource Isolation

Last updated : 2025-01-16 17:56:09

This document describes how Tencent Cloud TCHouse-D achieves resource isolation by configuring resource groups to avoid resource preemption and mutual interference among multiple businesses.

Note:

Only 2.0 and 2.1 kernel versions support configuring resource groups.

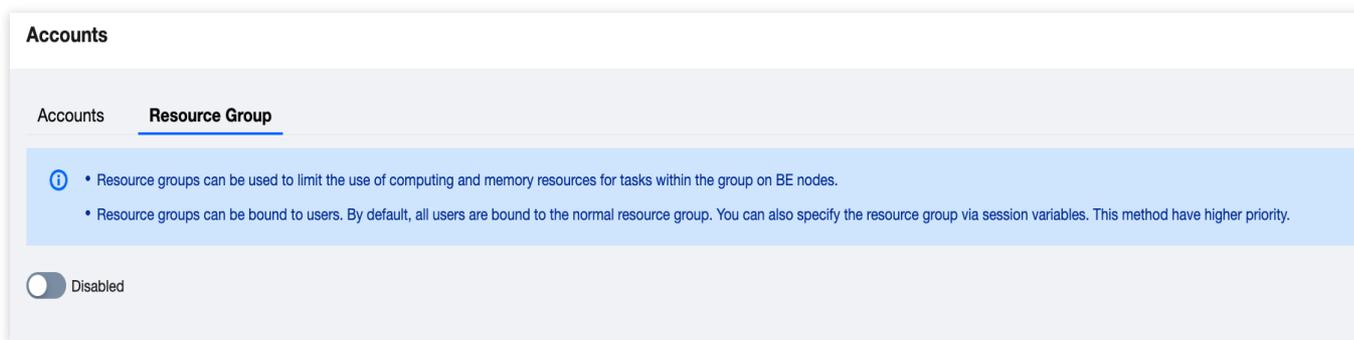
The 2.0 kernel version supports CPU soft isolation + memory hard isolation, while the 2.1 kernel version supports CPU hard isolation + memory hard isolation.

Operation Guide

1. Enter [Cluster List](#) and click **Cluster ID/Name** you want to back up.
2. Go to the Account Management page, click **Resource Group**, and enable the resource group configuration through the switch.

Note:

Enabling/disabling the resource group will synchronously modify the related cluster configuration, which is expected to take 2 to 5 minutes and will not affect your business during the process.



3. After enabling the resource group, you can click the **Add Resource Group** button to complete the creation of the resource group.

Note:

Once resource groups are enabled, a normal resource group will be created by default, and all users will be bound to the normal resource group by default.

Field	Applicable Versions	Description
Resource group name	2.0/2.1	Customizable, starting with a letter, supports numbers, letters, hyphens, and underscores with a length from 1 to 64.
CPU weight	2.0	Fill in for the 2.0 kernel version, the weight of CPU resources that this resource

		<p>group can obtain in each BE node (soft Isolation).</p> <p>Example 1: Three resource groups rg-a, rg-b, and rg-c are created with weights of 10, 30, and 40, respectively. At a certain time, rg-a and rg-b are running tasks, while rg-c has no tasks. At this time, rg-a can obtain $(10 / (10 + 30)) = 25\%$ of CPU resources, and rg-b can obtain 75% of CPU resources.</p> <p>Example 2: If only one resource group is running, regardless of its weight, it can obtain all the CPU resources.</p>
CPU usage (%)	2.1	Filled in for the 2.1 kernel version, the maximum CPU usage that this resource group can use in each BE node (CPU hard isolation between different resource groups).
Memory usage (%)	2.1	Filled in for both 2.0 and 2.1 kernel versions, the maximum memory proportion that this resource group can use in each BE node.
Memory soft isolation	2.0/2.1	<p>If memory soft isolation is enabled and memory resources are available in the system, the resource group can temporarily use system memory after its memory usage exceeds the limit.</p> <p>If memory soft isolation is disabled, tasks with the highest memory usage are canceled immediately once memory usage exceeds the limit.</p>

Create Resource Group ✕

Resource Group Name *

Starts with a letter, supports digits, letters, hyphens, and underscores with a length of 1-64.

CPU Usage (%)

The total memory usage of multiple resource groups cannot exceed 100%.

Memory Usage (%)

The total memory value of multiple resource groups cannot exceed 100%.

Memory Soft Isolation Enable

4. After the resource group is created, click the **Account Management** tab to bind the resource group to a user. Click **Actions > More > Modify Default Resource Group** to complete the binding of the user and the resource group.

Modify Default Resource Group



Username

admin

Bind Resource Group *

normal

Confirm

Cancel